

Engineering psychology: contribution to system safety

by Jan Noyes and Neville Stanton

The application of psychology in the design of systems is not new, and began in earnest after the Second World War with the formation of the Human Factors Society in the USA and the Ergonomics Society in the UK. What has changed very recently is the growing interest in the area of engineering psychology, as exemplified by a number of new developments. These include the 1st International Conference on Engineering Psychology and Cognitive Ergonomics held at Stratford-upon-Avon, UK, in October 1996, the formation of a Special Interest Group on Engineering Psychology within the British Psychological Society, and the launch of a new US journal in 1997—the *International Journal of Cognitive Ergonomics*. This article considers some of the major accidents which have occurred in recent years, and the contribution which engineering psychology makes to designing systems and enhancing safety.

In 1979, one of the most serious nuclear accidents occurred at Three Mile Island in the USA. For two hours and 18 minutes, the flow of coolant water to the main pumps was interrupted, and the plant operators realised that they were facing the possibility of a meltdown or at the very least an escape of radiation into the environment. The emergency continued for over 16 hours; there was no loss of life, but a small amount of radioactive material was released into the atmosphere, and the cost to the operating companies and insurers was estimated to be in the region of \$1 billion.

One December night in 1984, a gas leak from a small pesticide plant devastated the city of Bhopal in India. The company, owned by a subsidiary of the Union Carbide Corporation, was manufacturing the chemical 'methyl isocyanate', and the initial cause of the accident was thought to be an influx of water into a methyl isocyanate storage tank. At least 2500 people were killed, and more than 200 000 were injured.

In March 1987, the passenger and freight ferry *Herald of Free Enterprise* was making a routine cross-channel journey from Zeebrugge to Dover. Twenty minutes later, the ferry capsized, sinking in less than two minutes; 150 passengers and 38 crew lost their lives.

The above were all major disasters and, like Challenger, Chernobyl, the King's Cross fire and others, they are infrequent, unusual and unlikely to occur in exactly the same way again. Paradoxically, the outcome of many catastrophic events is that the situation is then far safer. After the experience at Three Mile Island, the design and procedures at nuclear power plants were extensively reviewed, with the result that nuclear power in the USA is now far safer today than it was at the time of the accident. These well-documented major events tend to become headline news and are usually followed by national inquiries into their causes in order to make recommendations to avoid similar incidents happening in the future. Inevitably, the reasons why these accidents occurred has to be attributed somewhere along the line to the fallibility of humans, since the systems are designed, managed and maintained by us.

ENGINEERING PSYCHOLOGY

Common causes?

The reasonable man [and woman] adapts himself [and herself] to the world: the unreasonable one persists in trying to adapt the world to himself [and herself]. Therefore all progress depends upon the unreasonable man [and woman].

George Bernard Shaw

Humans are creative: they attempt to make sense of the world by making inferences, solving problems and drawing conclusions. This capability together with our ability to adapt partly explains why we have been so successful in our environment. But in achieving this, we also make errors. Most errors are inconsequential with no lasting effects, some become positive events in that we learn from the situation and thus avoid repeating the error and precipitating the accident (the so-called 'near miss') and some have disastrous consequences.

Although it is generally recognised that human error cannot be prevented, and some would argue that it would not be desirable to do so,¹ we do know that the systems we use need to be well designed. Poor designs, for example, often do not take account of basic principles such as visibility, providing information about the affordances, possibilities and limitations of operation, and utilising natural mappings.

During the Three Mile Island incident, the operators failed to recognise that the pilot-operated relief valve was jammed open for more than two hours; the resulting water loss subsequently damaged the reactor, as it was supposed to open, relieve the pressure and then close automatically. In the Nuclear Regulatory Commission (NRC) accident report, the failure to diagnose this was explained as being due to the poorly designed control panel containing hundreds of alarms not organised in a logical manner. At the start of the problem, more than 100 alarms were activated with no means of suppressing the unimportant ones, and operators had to scan 1600 windows and gauges. Not only were the operators subject to an overload of information, several of the instruments went off scale making them impossible to read. This desperate situation was exacerbated by a printer failure, which resulted in loss of data and, when restored, a hard copy printout running more than two hours behind events.

The NRC accident report included the following:

(1) operators did commit a number of errors which certainly had a contributory if not causal influence in the events of the accident; and (2) these errors resulted from grossly inadequate control room design, procedures, and training rather than from deficiencies on the part of the operators.²


The Three Mile Island incident is not unusual in this respect: an analysis of other major incidents would indicate similar explanations. For example, Reason's comprehensive review and documentation of Bhopal, Challenger, Chernobyl, Zeebrugge and King's Cross³

indicated contributing conditions and latent failures resulting from poor design, insufficient training, inadequate procedures, system design (hardware and software) and regulatory, management, and operational failures, as well as communication and maintenance problems. However, there is always the difficulty when considering

latent failures of deciding how far back in the time frame to consider. This is important given that a catastrophic event arises from the coming together of several causal chains, where the elimination of any one would have deflected the sequence of events leading to the incident.³

In terms of demonstrating this, it may be useful to consider Zeebrugge, which could be described as an example of an accident waiting to happen, i.e. there were a number of resident pathogens already in place before that fateful day when the ferry set sail. For example, it could be suggested that the precursors of the accident started years ago with the design of the 'roll-on/roll-off' ferries themselves. These were designed for fast turn-around times, maximum profit and to load and unload vehicles as quickly as possible: bow doors open, cars drive in, doors close, ferry departs. Subsequently, their design was essentially a huge, great hull, which only needed a foot of water across the deck and the whole structure would turn over. Despite this, the design was accepted by the industry and roll-on/roll-off ferries became operational.

Admittedly, there were some concerns voiced at the time, but other designs were rejected primarily on the basis of cost. Further, a few weeks before the Zeebrugge accident, it was actually suggested to the Board of Directors that a light should be installed to indicate whether the bow doors were open or closed. This proposal was rejected, again on a cost basis. The owners had therefore made two decisions: one, to accept the roll-



Nuclear power station at Three Mile Island, USA

on/roll-off design, and two, not to fit the warning lights.

There were also problems with the management of the ferries: people were working double shifts, the boats were undermanned, and seemingly driven by profit-making over and above safety. The ferries had built up a culture of job demarcation; it was subsequently found that the person who was supposed to make sure the bow doors were closed was asleep in his bunk when the ferry left harbour, having just finished working a double shift. However, this was not an issue at the time of the accident—the ferries often left harbour with the doors still open. The crew would frequently close them as the ferry was leaving, and they had done this on many occasions without consequences.

On that fateful day in March 1987, the seas were choppy, the bow doors were open and water started to flood in. One deck hand noticed, but took the line that it was not their job, but someone else's. The layers of resident pathogens (i.e. the roll-on/roll-off design), poor safety culture (i.e. the negative reporting system), and poor management procedures (i.e. leaving port with the doors open) had effectively produced an operation 'riddled with holes'. Reason described this in his 'dynamics of accident causation' model, where a number of latent and active failures come together to produce the 'impossible' accident. These failures are indicated by the 'holes' on the model shown in Fig. 1. Although it might seem unlikely that a number of events would occur in order to create an accident (demonstrated by the holes lining up in the model) this can happen, as the Zeebrugge accident indicated.

A further point concerns who gets the blame. Typically, the people who are prosecuted are the individuals who committed the active errors on the day of the disaster. We tend not to look back in time to the people who designed the ferry, and the work procedures. The last person to be involved with the system tends to be the one who is blamed; in the Zeebrugge example, this was the person who failed to close the bow doors quickly enough. So what is the outcome? Often what happens is these individuals are taken to court, or prevented from having contact with operating the system in the belief that all is safe again. But it is not. The holes are still there, and undoubtedly people will make the errors again.

Contribution of psychology

Analyses of the cause of each of the cases mentioned above would indicate that the accident arose from a number of failures. These came from any one of a number of situations ranging from the design of the equipment, the work system, operating procedures, selection and

training of personnel, to the culture of the organisation, and human error resulting from poor decision-making at either a managerial and/or a lower level. It could therefore be concluded that accidents, and probably incidents, are multi-causal. However, this multi-faceted nature of the situation makes prediction of incident and accidents difficult, if not impossible. What then is the way forward?

All of the aforementioned systems involved people, and with even the most advanced technologies, humans have some degree of contact with the system. This may be during development and/or at the operational stage, if only to retain control of pressing the abort button in highly automated systems. Therefore, it would seem imperative to consider the role of the human, their capabilities and limitations, and characteristics in terms of psychological, physiological and social needs. Given the complexity of work situations, it is not only the actual interface between the user and the equipment which needs to be considered, but also the design of the workplace, operational procedures, the working environment and the culture of the organisation. This extends to encompass communication between individuals, their training needs, and the selection and recruitment of personnel.

Given the complexity of designing systems, there is likely to be a range of solutions from the retrospectively obvious to the diverse. A starting point might be to consider human error. This is an area which has been well researched by psychologists, although we are still some way from having a real, in-depth understanding of human error, and even the development of a precise definition continues to prove to be a difficult and elusive goal. As errors can arise because of a person's judgment in a specific situation, they cannot be defined objectively by considering the performance of humans or equipment in isolation. Furthermore, errors are caused not only by changes in performance with respect to the norm, but also by changes in the criteria of judgments. If we consider human error as occurring when 'any member

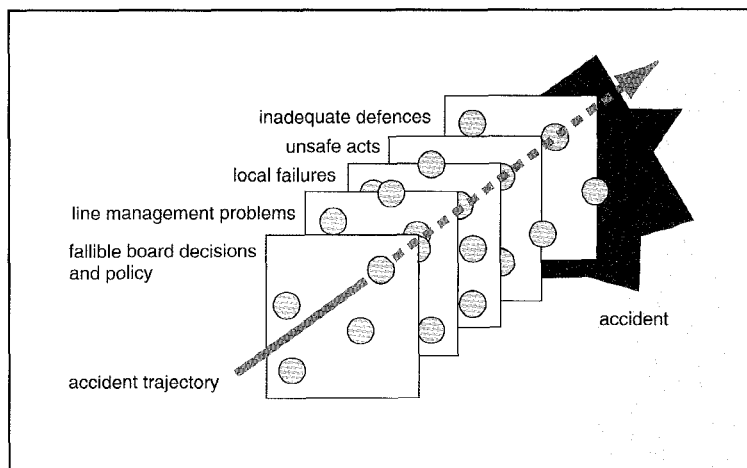



Fig. 1 'Swiss cheese' accident model, based on Reason's 'dynamics of accident causation' model³

ENGINEERING PSYCHOLOGY

*Herald of Free
Enterprise,
Zeebrugge, Belgium*



Rex Features Limited

of a set of human actions exceeds some limit of acceptability', it is possible to view it as an out-of-tolerance action, where the limits are defined by the particular system. However, human responses are subject to fluctuations or variations with respect to speed and accuracy. Kjellen⁴ put forward the hypothesis that an error or incident arises when the fluctuations exceed a boundary set by the working conditions. According to this model, the probability of a human error is dependent on the distribution of the fluctuation and the tolerance limits of the system. In this sense, errors are lawful in that they are predictable.

Complex cognitive tasks require optimisation on several criteria, and are subject to time-variant characteristics. Often there is no right or wrong decision, only the best decision for any given set of characteristics at a specific point in time. The outcome in terms of successful decision-making and avoidance of errors is often not known until later. When considering methodologies for studying error, this is very different from studying behaviour based on simple, rule-directed decisions. Laboratory studies of human error and post hoc analyses of incidents which involve collecting individual accounts/reactions etc. are often not fruitful in terms of yielding definitive information about the causes of making errors. Consequently, the most frequently-used methods for studying errors are observation and self-report techniques, and analyses of accident data.⁵

In summary, the complexities of human behaviour make studying human error a challenging task with many difficult theoretical and methodological problems. There currently exists no single theory or model for predicting the occurrence of human errors, and which would provide an initial step towards learning more about the causes and the prevention of errors. Consequently, when designing systems, failure to know in

detail why a human error occurs makes the development of a solution strategy both difficult and inefficient. In some safety-critical applications, there are also specific difficulties associated with studying human error due to the catastrophic nature of an accident when it occurs. As a result, evidence about the cause(s) of the accident may be lost and the main participants may be deceased, thus hampering the search for the causes of the errors.

Designing for error

One common viewpoint is that human errors arise because of a mismatch between the human and the task or human-machine misfits.¹ Frequent misfits are likely to be considered design errors, while occasional misfits may arise due to variability on the part of the system (component failures) or the human (human errors). Both external and internal factors may be responsible for this mismatch, although it is generally thought that internal traits, e.g. skill levels, are not as influential as external factors in their contribution to human error. External performance shaping factors of relevance here might include: (i) inadequate human engineering design, e.g. violation of population stereotypes resulting in sequence and selection errors; (ii) inadequate work space and work layout, which may contribute towards fatigue, decreased productivity and increased errors; and (iii) inadequate job aids, e.g. poorly written manuals and procedures, which may lead to uncertainty and errors on the part of the operator.

There are various strategies which can be employed to try to reduce the opportunities for human error, both within the finer details of the system design and those factors external to the system but directly relevant to its smooth functioning. In the first instance, the system should be designed in order to minimise the likelihood of errors occurring. When they do arise, it is important that

they are reduced to a tolerable level in order to lessen their impact on the system. Consequently, the recommendation is that errors should be reduced or eliminated through early detection.

With existing systems, an ergonomic audit involving a detailed analysis of procedures, operations and equipment design, in conjunction with studying the capabilities of intended users, may be needed in order to find the 'weaknesses', where human error is liable to occur. In addition, it is often useful for workers and experts to examine the work cycle in order to identify situations and processes likely to create errors. Outside the immediate operation of the system, there are various job supporting activities which might lead indirectly to a reduction in error rate. Continued training, refresher courses and simulator tasks may contribute towards reducing errors as workers are likely to gain from having time to reflect on new situations and review their performance.

Many human errors arise because of incompatibility between equipment design, procedures/training and achievable human performance. It is reasonable to assume that, as long as there is a human in the system cycle, the elimination of human error-related accidents is an impossible and unobtainable goal. A general design principle might be therefore to make systems error-tolerant.⁹ Interfaces should also be designed taking into account ergonomic principles. Some examples might include:

- The system and its components should be consistent and compatible with the users' expectations and mental models; the state of the system should always be unambiguously clear to the operator in order to reduce the occurrence of mode errors. Errors may be avoided through the careful and informed design of controls, displays and operational procedures.
- The system response should be reversible and observ-


able to the operators. Making the error more evident can be advantageous, e.g. programming errors in the lateral navigation of the Boeing 757/767 aircraft are shown instantly. If it is not possible to make actions reversible, they should be made 'difficult' to carry out in order to prevent unintentional performance.

- Systems should be designed to experience graceful degradation instead of total failure in response to a critical human error—an approach which has been used by NASA for many years. In this situation, back-ups or 'hot spares' (where equipment is kept running in case it is needed) are always available.
- Redundancy should be built into critical operations. Human performance can be monitored in order to verify that tasks are carried out correctly and at the appropriate times. It is rare (but not impossible) that a single error will precipitate a fatal incident. A lot of checks, redundant actions, reinforced with safety rules will ensure that several serious errors would have to be made in succession for a disaster to occur.

The ability of humans to adapt to situations has already been established as a positive human attribute; however, one of the drawbacks of this human trait is that we will adapt to poor designs. Equipment alone cannot protect against all potential types of human error; selection and training, procedures, support documentation, good communications and management plus operator self-knowledge are essential considerations in an error-tolerant system. When considering ways of alleviating human error, the recommendation is that the whole system including all the peripheral accompanying aspects is studied.

Conclusions

Accidents are usually multi-causal, and the resident pathogens in the design and operation of human-machine systems can lead to devastating consequences not only



Emergency services at the King's Cross fire



Chemical plant, Bhopal, India

for the workers themselves but also for people in the surrounding communities. Specifically, in each of the accidents mentioned above, operators were unaware of the seriousness of the system malfunctions because warning displays were poorly designed or located, and operators had not been sufficiently trained in dealing with these emergency situations.

The end result was that the Chernobyl nuclear disaster caused many deaths and exposed thousands of people to deadly radiation. The Bhopal tragedy, recognised as the worst industrial crisis in history, resulted in thousands of deaths and injuries to over one-quarter of a million people living around the chemical plant. One hundred and eighty-eight people died when the *Herald of Free Enterprise* capsized.

Most would regard it as unfortunate that it took disasters such as those mentioned above to force governments and organisations to consider the psychology more fully in the design and operation of complex human-machine systems. As an example, before the Three Mile Island incident, there were virtually no psychologists on the staff of the NRC or in the firms hired to design, build and operate nuclear power plants. Subsequently, and as a result of Three Mile Island, the NRC now has dozens of psychologists on its staff.

Since the 1940s machines and equipment have become more complex in nearly every industry. This, coupled with the continuing need to produce effective and safe systems, has resulted in psychology professionals being called to assist in designing even more efficient operating systems. In earlier times, a worker who made a mistake might spoil a piece of work or waste some time. Today, however, a worker's erroneous action can lead to dire consequences.

A final thought

If confronted with the operating panels of a nuclear power plant, or seated at the flight deck controls of a transport jet, our inability to perform would be of no surprise or consequence. In contrast, we would not expect

to be confused or make mistakes using the everyday objects which surround us: doors, telephones, video recorders, light switches, oven controls—the list is endless. But, how often do we try to push doors which need to be pulled? And when faced with a row of identical toggle switches, how often do we turn off the required light? When reprogramming the bedside clock radio, how many of us are faced with a bank of identical controls with no obvious clues to their different functions?

We all experience frequent difficulties manipulating objects in our everyday life, but is there any commonality with what happens in this context and in major disasters?

Prof. Donald Norman, Director of the Institute for Cognitive Science at the University of California, has written a number of books on the design of everyday objects, and in one he considers door design.⁷ Norman argues that a door poses only two essential questions: in which direction does it move? on which side should one work it? And when faced with a door, the answers to these two questions should be immediately obvious, with no need for words or symbols, or trial and error.

Well-designed objects are 'visible'; they have cues which make them easy to use, while poorly-designed objects lack these cues or give out information which is misleading. In fact, Norman suggests that doors which need to be labelled with their mode of operation are invariably poorly designed for human use. A trivial example perhaps, for what does it tell us about the design of more complex and advanced technological systems?

Well, there is one message it conveys—if we cannot get the design of a public door right for human use, what hope for the operation of advanced technologies?

References

- 1 RASMUSSEN, J.: 'The definition of human error and a taxonomy for technical system design', in RASMUSSEN, J., DUNCAN, K. and LEPLAT, J. (Eds.), 'New technology and human error' (John Wiley, Chichester, 1987)
- 2 MALONE, T. B., KIRKPATRICK, M., MALLORY, K., EIJE, D., JOHNSON, J. H., and WALKER, R. W.: 'Human factors evaluation of control room design and operator performance at Three Mile Island 2'. Report prepared for the Nuclear Regulatory Commission by the Essex Corporation, Fairfax, VA, 1980
- 3 REASON, J. T.: 'Human error' (Cambridge University Press, 1990)
- 4 KJELLEN, U.: 'Deviations and the feedback control of accidents', in RASMUSSEN, J., DUNCAN, K., and LEPLAT, J. (Eds.), 'New technology and human error', (John Wiley, Chichester, 1987)
- 5 NAGEL, D. C.: 'Human error in aviation operations', in WIENER, E. L., and NAGEL, D. C. (Eds.), 'Human factors in aviation' (Academic Press, San Diego, 1988)
- 6 BILLINGS, C. E.: 'Aviation automation: the search for a human-centred approach' (LEA, New Jersey, 1997)
- 7 NORMAN, D. A.: 'The psychology of everyday things' (Basic Books, New York, 1988)

© IEE: 1997

Dr. Jan Noyes is with the Department of Psychology, University of Bristol, 8 Woodland Road, Bristol BS8 1TN, UK, Tel: 0117-928 8560, Fax: 0117-928 8588, E-mail: J.Noyes@bristol.ac.uk. Neville Stanton is with the Department of Psychology, University of Southampton.