

Chapter 7

Computer-Centered Computing: What Are “Human Factors” and Why Should We Care?



Now where's that warranty?

Abstract Have you heard the old joke about the fellow who finds another chap searching for his lost car key under a single street lamp on a dark night? He offers to help and gets down on his hands and knees on one edge of the circle of light, but after a half an hour feels like he's not having any luck.

“Are you sure you dropped it here?” he asks and the other fellow answers:

“Here? I didn't drop it here. I dropped it about a hundred and thirty paces back, over on the other side of the street.”

Upset at having wasted his time, our hero asks why they’ve been searching on the wrong side of the road, and gets the following reply.

“Don’t be stupid! There’s no light over there. We’d never find my key in the dark.”

With that in mind, let’s consider a simple task, say, reading a sentence in a book.

I’m going to presume that you’ve just read that last sentence, and use it for my example. Now, if we were to try and express the task of reading that last sentence mathematically, there would be many factors involved. To reduce them to a useable example, and for the sake of expediency, we’ll ignore all of the factors that went into how you came to be reading that sentence, and just work from there.

Could you see the book and the words in it well enough to read them? Is there enough light, or maybe too much? Are the letters expressed clearly enough to be read, or do they seem blurry and out of focus? Should they be more sharply defined, or presented in greater contrast to their background? Maybe they should just be bigger? Maybe they should be written in a different language, presented in a different medium, or available at a different time or place.

The Human Factors


In mathematics, the factors of a number are all of the other numbers that act together to make it up. In this way, 1, 2, 3, and 4 are one set of factors of 24, because $24 = 1 \times 2 \times 3 \times 4$.

Mathematical factors are pretty straightforward. What are the factors of taking a walk, or of driving a car? As a first approach, we can see that some of the factors must be environmental (weather and geomorphology will have an effect on one’s ability to drive or to walk) and some must be technological (the nature of a built path or road is a strong factor in one’s ability to follow it, as are the engine, the surface of the tires, and the fit of one’s shoes). These factors are all in the domain of specialist engineers and technicians.

But we must also allow that some of the factors are human. How well can the traveler see? Has the traveler learned to drive – or to walk, for that matter? It becomes more complex when we factor in emotional states, as Einstein reasoned in his famous illustration of the relative nature of time. Adding more people to the equation means considering their individual skills, values and emotions, and also demands a consideration of their interactions.

It is easier for engineers to quantify technical and environmental factors than human factors. Because of this, human factors are considered a peripheral specialist domain and are largely ignored by other engineers during the development of either simple or complex systems. Despite their relegation to the sidelines during the design and development stages, they are often the focus of investigations when a complex system fails. Why do such investigations routinely blame individuals? For that matter, if the failure was due to human error, was it in the use of the technology, or in the design of it?

Over the next three chapters we'll look at this general concept of human factors as opposed to technological and environmental factors. In Chaps. 8 and 9 we will delve more deeply into the applied biomechanics of physical ergonomics and the applied psychophysiology and neurology of psychological ergonomics, but first, let's spend a chapter talking about how Human Factors are applied.



Look, I'm not denying that it's getting warmer, I'm just saying I don't believe that it's caused by frogs

What are the Human Factors that are shaping your experience of reading this chapter? Do you need corrective lenses, or the chance to focus with one eye at a time? Maybe you need the text to be further away from your eyes, or closer?

Do you recognize the alphabet? Do you recognize the words? Can you attach meaning to them at all, and if so, can you do so with them in this combination?

Maybe you are too hot or too cold to focus on reading, or there is wind blowing in your eyes and blurring your vision? Maybe the words are bouncing and you cannot focus on them and trying has upset your stomach? Maybe you and the words are both bouncing in unison, and trying to read has strained your neck muscles?

Maybe there is something in your culture or in your personal understanding of the world that prevents you from reading here and now.

Maybe there's a distracting noise in the background, or the person beside you keeps farting and you don't know whether to be disgusted or laugh out loud. Maybe the person behind you keeps kicking your chair.

Maybe you're too tired to read, or something has caused your eyes to be less functional? Maybe you could read but you are overcome by anger or sadness and

cannot focus? Maybe your lack of focus is based on boredom – this example is getting to be rather long-winded. Maybe that is an excuse you use to hide the fact that you have emotional defects that prevent you from reading, or maybe you are drunk, or drugged, or just plain stupid.

Maybe you feel that this book has just insulted you, and now you don’t want to read it anymore.

Any and all of these factors could affect your ability to perform the task of reading this chapter. In fact, the best-designed machine in the world for text display, running the best software in the world for text display, on a customized operating system that is working at 100 % efficiency could still fail to provide you with text that you can read, due to any one, or any combination, of the human-centered reasons listed above.

Those are the human factors, and they are largely ignored by the engineers, technologists, and computer scientists who have built most of the tools we use every day: machines and tools made up of components they have studied, for use by humans, whose components they have not studied at all. Unfortunately, they prefer to do their work in the circle of light they know, rather than venturing into the darkness to find the key.

Even that wouldn’t be so bad, if that were the end of it: I begrudge no one their fears. But even if they want to confine themselves to the circles of light where they feel at home, they should allow some of us to explore the darkness.

Humane Factors



Officer, there is no doubt whatsoever that the cause of death is human error!
My report clearly shows that there is nothing wrong with this rock!

Back in the first decade of this millennium, I was hired as the Human Factors specialist for Safety Intelligence in the civil aviation branch of Canada's Ministry of Transportation. One of the first things that happened to me after settling into my office in Ottawa is that I was sent to Las Vegas to attend a course that had been arranged by the woman I'd replaced.

The course was an introduction to a very popular method of accident investigation. The creators of the method were teaching the course and certifying the participants in the use of their very popular software. I won't mention the names of the people or products, because the story I'm about to tell does not shine a nice light on any of us. Like the other stories in this book, this one is true to the best of my recollection, but it is only my version of a series of events that happened a long time ago. If anyone is offended by this story, I ask that they please consider all of the parts they dislike to be figments of my poor memory. The story is told without malice or intent to harm. It is told as an example of how perspective often misleads human factors investigations.

So there we all were, 20 or 30 human factors specialists from around the world, sitting in neat rows and watching presentations being put on by 2 collegial gentlemen. The course went well, but I should admit before I go any further that I did not then, and do not now, subscribe to any of the theories that espouse simplification and categorization as necessary to accident investigation. The truth is that I believe that the system being presented that day is very good for finding scapegoats, and very bad for finding preventable contributing causes. In other words, I think their system is based on a false concept of how their work should be done.

That said, I think they've done it well. If your intent is to build a strong evidence-based case for blaming someone for an accident, then I could recommend their system and their software... so long as you don't care whether or not the individual you identify should actually be blamed for anything more than one in a series of mistakes, most of which were likely made far away from her in both space and time.

Accident investigation is supposed to be about finding out how an accident happened, with the goal of learning from it and generating remedial actions that will either prevent that particular type of accident from ever happening again, or further limit both the conditions under which it is unavoidable, and the impact it will have when it does occur. Too often, especially in private corporations and in the military, the purpose of accident investigation seems to be to find out who you can blame for what happened, so that political and financial liability can be managed with minimal impact on the entity that employs you.

The men running our course were former military who now made their living as consultants to private industry. So their perspective may well have been doubly-reinforced. All the same, I would not have disagreed with them and made a public scene among my peers, if not for the fact that they did something that I found particularly offensive during the course.

Towards the end of the course, one of the instructors offered up an example from their current caseload, and provided details about a fatal accident that was still under investigation at the time. They walked us through the entire event, even though that allowed a number of people in the class to loudly proclaim that they knew which

accident it was. That is a real professional “faut pas”, but it’s not why I spoke up. The instructors even had the gall to walk us through to a conclusion that was patently false, blaming the one person who died for the entire incident and absolving the two corporations that violated their own rules and procedures to make the accident possible. I ignored that, too.

The thing I couldn’t ignore was the message they gave us along with their pronouncement of the dead man’s guilt.

Let me tell you a condensed version of what they told us. An experienced mechanic had been gruesomely killed in front of a great many witnesses while trying to figure out why an engine wasn’t responding properly. The accident was, of course, his death, and they explained at great length how their system allowed them to pinpoint the personal problems that had caused the error in judgement that resulted in his death.

They didn’t address the fact that the whole thing happened in front of witnesses, except to offer brief commentary on the potentially catastrophic civil suit that could be put against both the airline and the airport. The location of the accident was the key to the whole thing.

You see, the airline, faced with having to run tests on the engine, should have emptied the plane of passengers and sent it off to a service hanger. They should have, but that would have seriously delayed their flight and required them to deal with the expenses of taking a plane off-line and finding a replacement – not just for that flight, but for all of the other flights that plane was due to make before its next servicing.

The airport, faced with an airliner that needs this kind of service is supposed to require the plane to move away from the passenger area, and into an area dedicated to testing and repairs, like the previously-mentioned service hanger. Insisting on that protocol being followed would have forced the airline to make the more expensive decision.

Instead, the decision was made to call in a contractor who could hurry through testing and repairs. It was this hurrying that made it necessary for a man near retirement to do the job, with only a new trainee to help him, and it was this hurrying that had the plane sitting next to an occupied plane and in front of the big windows of a boarding area, when the error was made and the man died.

If the plane had moved, no passengers would have seen the accident. More importantly, if the plane had moved, the immense time pressure would have been entirely removed, and the entire operation could have been carried out by the book. In fact, the man who died would probably never have touched the plane that killed him that day.

So when the instructors joked about how easy it was to isolate the person responsible, using their privately licensed software, I had to point out that the software was only making it easier for them to justify the assignment of undeserved blame to a victim, and the absolution of their corporate client.

I passed the test and got my certification, but the instructors and a number of my fellow students did not think kindly of me. Some of the other students told me privately that I was right...

...that I was stupid to think the system would ever change, but that I was right.

Maybe I was being stupid, and maybe I was just being optimistic. Certainly, I was still holding on to this notion that scientifically-minded people share their ideas in order to test them; that a scientist would welcome valid criticism of hypotheses and theories, and even of methods and conclusions.

The problem was that it's not just pilots and mechanics and conductors and drivers who make Human Factors-based errors in judgement.

Investigators do, too.

What Are Factors and Which Ones Are Human?

As mentioned earlier, in mathematics the factors of a number are all of the other numbers that act together to make it up. In this way, 1, 2, 3, and 4 are one set of factors of 24, because $24 = 1 \times 2 \times 3 \times 4$. In the same way 4 and 6 are another set of factors of 24, whether or not we include the number 1. In fact, the number 1 is a factor of every number because its presence has no effect on the other members of the set. We'll come back to this in a little while.

Aside from being an art unto itself, we can use mathematics to model ideas that reflect our understanding of some aspect of the real world, and so test those ideas theoretically. This is what an engineer is doing when she calculates the range of potential joint strain on the buttress of a bridge in order to determine the material factors that will allow it to be made and used safely. It is mathematics that tell us the gradient to use on the upper and lower surfaces of a plane's wing, so that the wing can pull the plane up into the sky. Applied mathematical models do a fine job of holding up bridges and planes and they hold up in many other areas as well. The only problem is that – by their nature – models are simplifications. We have to consider whether the factors in our model represent every important factor in the real-world system.

For many years, the aviation industry resisted the idea of even considering human factors in the modelling of flight. Captain and crew were expected “simply” to perform without making any errors [4]. I hope that we have reached a point in our general popular knowledge now, from which we can look back at that idea and recognize it as a ridiculous example of false pride. You see, humans make mistakes. Not some humans, but all humans. It is in our nature to have physical and psychological limitations, and to fail if we try to surpass them. It is also in our nature to fail, occasionally, even when we are performing a familiar task in a familiar setting.

I'm sorry, what's that? You say you know someone who never fails at familiar tasks? That it is just a question of developing a high level of skill with the tool in question, or that it is just a question of performing 10,000 h of practice? That's an interesting theory. When was the last time this acquaintance bit their tongue? Was it due to a lack of familiarity or to a lack of practice?

This conscious acknowledgement of our propensity to fail is crucial to designing and executing tasks so that they can be completed safely. We must try to understand

the human factors that may contribute to the safe completion of the task at hand, or to the failure.

Let’s return to the examples of taking a walk or driving a car. As mentioned earlier, we can immediately see that some of the factors must be environmental. For example, 1 m of snow on the ground would have an effect on one’s ability to drive or to walk, as would a steep downwards gradient. Along similar lines, the surface texture of a built path or road is a strong factor in one’s ability to follow it, and so are other technological factors. One might consider every aspect of the tool being used, whether it is a car, a truck, or a pair of boots. Each of these factors is the domain of specialist engineers and technicians working for corporations, whose concern is the steady, marketable impression of improvement over time.

These factors contribute to the experience of walking or of driving in very fundamental ways but, can we really consider the either experience without giving some thought to the person or persons who will be doing it?

Some human factors are fundamentally physiological. How well can the traveler see? Do her legs function well enough for her to drive, or to walk at all? If so, has she learned to drive - or to walk, for that matter? The factors quickly become extremely diverse, without losing any importance. Is she mentally prepared to travel safely, or is she too agitated, too pre-occupied, or too tired. The complexity increases when we factor in one or more additional fellow travelers. Not only must we ask each of those earlier questions again, we must also consider the effect each person will have on every other person. Again, this is not a simple matter of saying the people “get along”.

How will each traveler affect the experience of the others individually, as sub-groups, and as an overall group?

There are well established metrics for measuring the friction of a road surface, and there are well-established standards for determining the range of safe and un-safe grades or slopes or turns. There are even international metrics and standards for measuring skill at driving, and for providing both static and dynamic safety information to travelers.

But what is the metric for measuring how well two people will get along? Does it change based on the type of task each is performing, or based on whether the two are in a larger group?

How does being tired affect your skill as a driver? How about the presence of two friends, does that affect your performance? Exactly how does the effect differ if the two of them love each other, or hate each other? These human factors are not clearly defined. This means that they do not have a direct measure, which means that they are not easily quantifiable, and that means that they are not easy to formally incorporate into any practical model of the factors that contribute to driving.

Are we ready to admit that our model of driving doesn’t represent enough of the real world problem? What about all of our other models?

What Is and What Could Be in Human Factors: What If Safety Were More Important than Profit?

We have discussed Maslow's Hammer and the idea of "déformation professionnelle", so let's apply them to the question of why Human Factors investigations routinely blame individuals.

Could it be that their tools measure human failure better than they measure other failures?

Even if the error was human, was the failure in the use of the technology when and where the incident happened, or was the human error made by the designer of the technology, or by the manager who sets or enforces policy about how it should be used?

Have you ever read any of Richard Feynman's autobiographies? In the first one, "Surely you're Joking, Mr Feynman", the Nobel laureate shares stories from different periods of his life, including his time on the committee investigating the cause of the explosion of the US Space shuttle "Challenger" in 1991 [1]. I won't spoil the story here, but the book is easy to find – just ask your local librarian. Yes, yes you can just search for it on-line instead... ..but I think you're missing a great opportunity if you pass up the chance to discuss science books with a knowledgeable librarian.

Now, if you hate the idea of hunting down a story that I could just summarize here, then... well, then you'll probably also hate my next suggestion, and the one after that, too. I suggest that you stop and reflect for a little while on all of this hatred. It used to be a much stronger word than it is now. I suggest that you put some thought into whether or not you could use the word "dislike" instead.

But I digress. I hate it when that happens. Let's take a look at one of the basic tools of the trade of Human Factors.

The idea of human factors seems to me to be fairly straightforward. I hope that, if you have read the earlier parts of this chapter, you might agree with that opinion. I find it interesting that, during my tenure as a Human Factors Specialist, I dealt with so many people who were convinced that the whole concept was hard to understand.

I mean, so near as I could tell, most of them were human. I think that puts you at a natural advantage... or at least that it should give you a good starting point.

Well, I did meet one person who clearly thought the field was even simpler than I do. She was the political appointee whose job was to decide if I would be allowed to publish or present any information to the public. Her HF training had been a one- or two-weekend seminar. What's more, she didn't feel the need to apply any of the information she might have gained during that training. Her job was just to enforce the policy of the government at the time, which was that the only science worth reporting is science that directly supports some government policy.

Really, that happened, in Canada, in the Twenty-First Century. I won't go into detail about it here and now, but will encourage you to look it up on your own. You




Fig. 7.1 The human-centered relationships reflected in the SHELL model

don't have to look for it among conspiracy theorists or anything like that. You can check the archives of any of the major Canadian news services.

What I will say about it before leaving it behind until the end of this book, is that the policy seems to me to neatly reflect a very normal form of human behaviour – the kind of self-delusion that we are able to use to convince ourselves that terrible, illogical decisions are actually rational and justifiable. I'll discuss that mechanism in Chap. 9. For now, let's discuss one of the tools that shape the way that Human Factors specialists try and solve their problems in my old field of aviation.

The SHELL model was first proposed by Edwards, in 1972 [2]. It was an attempt to simply represent the fundamental concepts of human factors so that others could apply them. In order to explain the model, we will refer to the modified SHELL model as it was illustrated by Hawkins in 1975 [3]. There are five images presented in Fig. 7.1. The first is the SHELL model itself with its five components. The next four images are intended to illustrate four possible relationships between the central component and the four on the periphery.

Let's go through them in turn.

The basic model shows five tiles with rough, irregular edges. The central tile is labelled “L” for “Liveware” or living creature. This is the human at the center of the task. That human has to interface with four general things. Software (S), Hardware (H), the Environment (E), and more Liveware (L) such as co-workers, management, clients, etc...

The borders of the tiles are rough and irregular in order to represent the rough and irregular interactions that take place in the real world. We see these relationships modelled in the four smaller images. Starting in the top row, next to the larger model, we see that the relationship between the central human and the software they must use. To the right we see the relationship that the user has with their hardware. Bottom left shows the relationship between human and the environment in which the work is being done, and the final image models the relationship between the

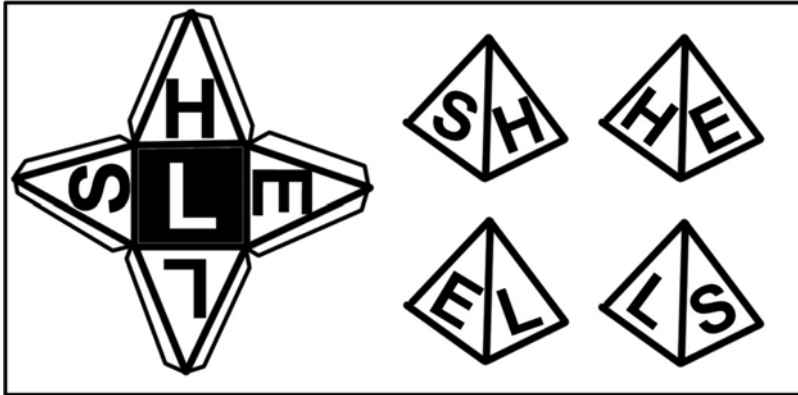


Fig. 7.2 My modification of the SHELL model, in which each individual factor is assumed to be human-centered, unless it is considered in one of the four other direct relationships

central human and other humans. Some have argued for the addition of a C for culture, but I believe that the way that L interacts with everything else is fundamentally based in their culture, so the C must be assumed to be a part of each L.

I have always found this model to be too simplistic. Each human on the job must have a similar network of connections, and so must each of the other tiles. Hardware and software must be compatible, so there should be some connection between the two. A poor or missing connection would be very important and it would be a clearly measurable and improvable problem. Just the kind of thing that a model should reveal. Hardware that suits some environments does not suit others, there should be a link there as well. When I was working in Human Factors for Civil Aviation, I proposed a series of more complex models that were intended to better reflect that depth of interrelationships.

My final attempt was a pyramid. The letters are the same, but the peripheral tiles are differently-shaped. If they are each triangular, points outwards from the center, then they could be folded up into a pyramid in which software is also connected to hardware, and hardware is also connected to the environment. Software and environment are both connected to the outside liveware (Fig. 7.2).

I was happy with that for a little while, because it illustrated all of the relationships intended in the original model, but also served to make all of the connections more clear. What's more, it moved the user from the center of the situation to being the basis of the model. I thought that nicely reflected the idea that user who is central to an investigation is not always the center of our focus but, as shown in the pyramid, should be the basis of our work. The truth is that the complex interrelationships between multiple people and multiple devices would be better represented with something more akin to a rhombic triacontahedron (30-sided object) in which the shape of all component surfaces reflected their function and relative importance to the group.

In the end, it proved too complex for me to make a single geometric model that reflected everything I wanted it to do. What’s more, it was overwhelming to start talking about geometry when trying to explain the importance of human factors. Most of the people I was working with hadn’t given a thought to theoretical geometry since high school. If they applied it practically (in terms of their work in flying or maintaining planes), they did so without thinking of it as geometry. I did carry the pyramid with me, though, and I’d hand it to the people with whom I was discussing problems, just so they could see that any issue with any element came back to the person, and any action by the person would be reflected in one or more of those elements.

It was fun, but I thought that there had to be a better way to talk about how humans interact with their environment and the other things in it. I wanted to show that categorizing the elements with which we interact was not the important thing. This always made up too big a part of discussions when people tried to apply the SHELL model.

XERES: “Well, does our problem with the air conditioning get sorted under E for environment or under H for hardware?”

JERRY: “You’re asking about fitting the problem with the air conditioning into a SHELL model?”

XERES: “Yes, yes I am. In fact the air conditioning is supposed to be monitored by a central computerized system, so maybe the problem should be under S.”

JERRY: “Actually, the real problem with the air conditioning is you, Xeres! You’re picking at the manual controls three or four times every day! So, yeah, we could file it under the second L, or you could just stop mucking about!”

Just so you know, the above is not an accurate transcript of a real conversation I heard during my days as a specialist, but there were far too many times that people were trying to make the facts suit the model, rather than trying to use the model to address the facts.

This led me to develop a very different model. Figure 7.3 is my model of General Human Interaction (GHI) in which one’s thoughts, one’s immediate environment, and the world at large are categorized as either restricted (inside a square) or unrestricted (inside a circle). The three differently-sized pairs of squares and circles were intended to represent one’s interaction with the world at large, with one’s work environment, and with the same kind of conflict between one’s own deliberate and task-minded thoughts, and one’s more emotional and impulsive reactions. While I found the model useful, it was hard to make it seem practical to people who were less concerned with why something went wrong than with fixing it. That level of practicality keeps planes taking off on time. Finding the underlying factors behind recurrent problems and learning to prevent them, that’s what keeps those same planes in the air.

I came back to this model in developing Anthropology-Based Computing. You’ll see that it looks quite different when we find it again in Chap. 10.

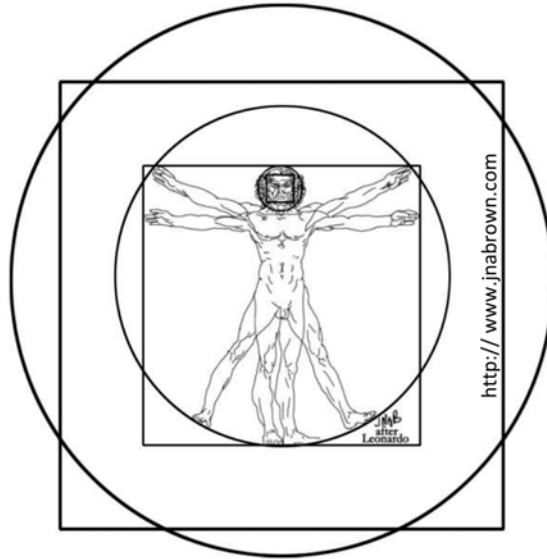


Fig. 7.3 A model of General Human Interaction (GHI) in which one ‘s interactions with the world at large, with one’s immediate environment, and even one’s thought are forced to adapt to externally-determined restrictions (inside a *square*) or are allowed to occur intuitively and in a manner that seems natural (inside a *circle*)

I want to leave you with an explanation of two standard approaches to investigating HF-related issues in aviation. James Reason [5, 6] explained the dichotomy of as follows. The human approach “...views these unsafe acts as arising primarily from aberrant mental processes such as forgetfulness, inattention, poor motivation, carelessness, negligence, and recklessness”, while a system approach, is based on the belief that “...though we cannot change the human condition, we can change the conditions under which humans work.”

I used to agree with that statement whole-heartedly, but now I have come to believe that changing work conditions to suit workers is a very basic first step. I think that the next step is to learn how to adapt our tools so that it becomes reflexive and intuitive for the worker to do their job accurately and without the risk of undiscovered failure.

To do so would require a better understanding of how the body and mind really, naturally, work.

Summary

Well, there *was* a craft you could make with scissors and glue, and I think the cartoons were funny...

References

1. Feynman RP, Leighton R, Hutchings E (1985) “Surely you’re joking, Mr. Feynman!”: adventures of a curious character. W.W. Norton, New York
2. Edwards E (1972) Man and machine: systems for safety In: Proceedings of the BALPA, Technical symposium, London
3. Hawkins F (1976) Some ergonomic aspects of cockpit panel design for airline aircraft. Shell Aviat News 437(438)
4. ICAO Circular 216-AN31 (also called Human factors Digest No 1), republished as CAP 719 in 2002
5. Reason J (1990) Human error. Cambridge University Press, Cambridge, UK
6. Reason J (2000) Human error: models and management. *BMJ* 320(7237):768–770