

DISTRIBUTED COGNITION AND HUMAN FACTORS FAILURES IN OPERATING AND DESIGN PROCESSES

J S Busby and E J Hughes
Department of Mechanical Engineering, University of Bath, Bath BA7 2AY

J V Sharp and J E Strutt
School of Industrial and Manufacturing Science, Cranfield University, Bedford MK43 0AL

E Terry
Sauf Consulting Ltd, 30 Observatory Road, London SW14 7QD

The principle of distributed cognition provides a promising framework for understanding the way that human problem solvers rely on the environment to accomplish their tasks. People rely on cues they receive from human co-workers, they make inferences from the appearance of the artefacts they work with, and they draw on organisational culture to work out what is expected of them. This distribution of cognition is vulnerable to particular kinds of failure, however, and such failure can occur both in the design of an installation and its subsequent operation. The purpose of this project is to use the distributed cognition principle *both* to help the designers of offshore installations reason about human factors connected with their designs *and* to help the same designers reason about how they should work effectively in collaboration with human factors specialists. We also wish to find out how such aspects of the designers' environment as safety regulation and safety management systems influence this distributed cognition. The aim is to develop two main products: a workbook for offshore design organisations to help them anticipate human factors failures, and a guidebook for regulatory organisations and senior industrialists on how the safety environment influences failures and limitations in designers' thinking.

Keywords: human factors, human error, engineering design

INTRODUCTION

In one of our recent studies[1], we analysed a case in which a maintenance engineer was killed when he used a beam inside a vessel in order to suspend a hoist. The designers might have forestalled the accident if they had been able to predict this kind of behaviour at the accident site. They might also have forestalled the accident if they had realised there were certain cues that indicated they should consult human factors specialists during the engineering process. Furthermore, it is possible that with a more influential safety management system, a more influential professional culture and more influential regulation the designers would have actively sought such cues instead of responding only to those brought to their attention. The lessons of accidents such as this are therefore that:

- designers need a model to help them reason about failures in the process of operating the equipment they design;
- design organisations need a model to help them reason about failures in their engineering processes;
- regulators, professional institutions and managers need an understanding of how the regulatory, professional and managerial environment influences designers' reasoning about failures.

The principle of distributed cognition provides an organising principle that, potentially, can provide a consistent way of tackling these needs. This principle is that an individual's learning and problem solving is explained in terms of his or her environment and the tools, people and procedures that the individual works with. What you do and what you learn do not arise from mental processing alone but mental processing in conjunction with this environment. Similarly, when you fail the failure is not simply a limitation of human information processing but a joint limitation of human information processing and the features of the environment at the time. If, for instance, a fitter installs a check valve the wrong way round we should look at what characteristics of the valve, what features of the fitter's particular experience, and what assumptions made in the prevailing culture gave rise to the fitter's erroneous model of the valve's required orientation.

There are both weak and strong views of distributed cognition[2] - the weak view saying that there are both 'solo' and distributed cognitions, the strong view that all cognitions are distributed in some way or other. In both views, however, expertise and intelligent behaviour is seen as a characteristic not of individuals but individuals interacting with a technology[3, 4]. Applications of this idea have been mainly in real-time tasks where the environment provides much of the pacing and the human participants have little time to reflect on their current actions. Recent examples include, for instance, air traffic control[5]. Our own recent work has tried to apply the principle to understanding design processes, both where they succeed[6] and where they fail[1]. Distributed cognition is related to the ideas of situated cognition[7] and external cognition[8] - both of which again emphasise the role of a person's environment in that person's beliefs about the world and problem solving within it.

The idea in this project has been to apply this principle to understand failure in two processes - operating a piece of equipment or system, and designing it. Ultimately, we want designers to reason effectively about both how they contribute to operating failures and how the design process itself can fail. Also, we want them to reason about both kinds of failure in the same basic way because we want them to understand operators' failures as being of the same kind as their own failures. Designers sometimes make mistakes when they are under time pressure and take short cuts with CAD tools, and typically blame the CAD tool for providing too little feedback on the consequences of certain actions (like turning layers off). They need to see the operator of the plant they are designing also as a person under time pressure looking for short cuts in the operation of the plant who will make mistakes when the plant provides too little feedback on the consequences of the operator’s actions.

Figure 1 shows the scope of the project. It shows distributed cognition in both design and operating processes, and shows how we would like to model this distributed cognition and the way the environment influences it.

OBJECTIVES

The gist of the project is to analyse past failures, use these to develop distributed cognition models of failure, and build these models into practical tools. In more detail, the specific objectives have been as follows:

- To investigate how distributed cognition failures contribute to accidents in the process of operating offshore installations.
- To investigate how distributed cognition failures contribute to shortcomings in the process of designing offshore installations.

- To investigate how the regulatory, professional and organisational environment influences these failures.
- To develop models and workbooks that guide and support the design organisation.
- To develop models and guidebooks for the Health and Safety Executive (HSE) that help staff involved in inspection, audit and policy development.

Our intention is that there should be a number of benefits:

- One of the main difficulties that designers have in thinking about human factors is their lack of a general model that helps them enumerate potential problems systematically. This work ought to provide such a model.
- Similarly, one of the main obstacles in all collaborative work - but particularly collaboration involving a technical discipline and a social discipline - is the lack of a model that helps each discipline ask the right questions and send the right cues to the other discipline. Again this work ought to provide such a model.
- Because accidents are relatively infrequent, it is very hard to be confident that one organisation's historical experience is enough to help it predict and prevent future accidents. Having a model gives some assurance that there are not large gaps in an organisation's understanding of how things can go wrong.
- A model of this kind would also help reveal any limitations and gaps in current HSE guidance on incorporating human factors in the offshore design process. In particular, by understanding both the manner in which cognition is distributed, and how this is affected by such factors as regulation and knowledge of regulation, we hope that it will become clearer how to influence designers' thinking.

METHOD

INVESTIGATING FAILURES IN THE OPERATING PROCESS

This has involved three main steps:

1. Building a case base of past accidents and incidents.
2. Identifying the distributed cognition that has failed through the analysis of these cases.
3. Developing classifications of these failures.

The cases have been obtained from two main datasets. The first is a fairly diverse collection of public domain investigative reports, particularly the Marine Accident Investigation Branch (MAIB) digests, the HSE's UK Continental Shelf Risk Review and proceedings of public enquiries into offshore disasters. The second dataset is the HSE's incident and early day report database. The analysis of each case has involved asking two main questions:

1. In what ways was information processing distributed? (For example, did design and installation engineers both have to have certain bits of consistent knowledge in order for a valve to be installed correctly?)
2. In what ways did this distribution fail? (For example, did the installation engineer draw on a memory of a previous valve instead of examining the new valve when determining the correct orientation or fitting?)

The next step has been to look for general patterns of failure and develop a classification. This is the stage that we have reached at the time of writing, described briefly below in the Results section.

INVESTIGATING FAILURES IN THE DESIGN PROCESS

This part of the work involves failure in the design process, rather than accidents in the operating process, and at the time of writing is not yet underway. Failures of the design process, such as in the collaboration between designers and human factors specialists, are naturally harder to obtain because they do not usually reach the public domain. Our intention is to draw on three main sources of cases. The first is a study of error in the design process that we conducted recently with several design organisations. This yielded a database of about 100 cases. The second source is a set of experts from different disciplines, in our own institutions, who have been involved in consulting with firms on failures of various kinds. This includes a reliability engineer, a regulatory expert, an organisational sociologist and a psychologist specialising in ergonomics. The plan is to run elicitation exercises with them in order to draw as systematically as possible on their experience in this and related industries of failures, breakdowns and limitations in the design process which have introduced hazards in the equipment being designed. The third source is a group of engineering designers in an offshore installation design organisation. Our intention is to get designers' observations of how the design process has failed - especially in the collaboration among different disciplines, especially in the involvement of human factors specialists.

As with the accident case analysis, the plan is to identify in these cases the nature of the distributed cognition and the modes by which it failed - and build a classification of these.

INVESTIGATING THE ENVIRONMENTAL INFLUENCES

The third part of the work concerns external influences, especially the managerial and regulatory environment. This has not yet started.

RESULTS

The work has at the time of writing been underway for four months and, accordingly, only the first main element has been tackled: failures in distributed cognition during the operating process. In the rest of this section we have described some of the general patterns that have been identified.

RELIANCE ON CULTURAL ARTEFACTS

In one case of an offshore platform capsize, the designers had relied on existing compartmentalisation standards that seemed to have been inappropriate to a structure of this kind. This led them to overlook the possibility of partial capsizing (which eventually became a full capsize). One consequence of the failure to anticipate a partial capsize in a semi-stable condition was that the lifeboat davits would not allow a launch in this condition.

The designers' problem solving was effectively distributed since they were relying on partial solutions provided by engineering standards. This distribution failed because there was no apparent inspection of the standards to test their applicability to the idiosyncrasies of this application. More generally, standards are 'cultural artefacts': things developed historically that get passed down to future engineers. This case provided an example of how using cultural artefacts can lead to failure because the conditions in which they are developed are not usually identical with those in which they are subsequently applied. A culture, typically, does not keep up with a technology.

RELIANCE ON CULTURAL ASSUMPTIONS

Another contributor to the same accident was that the designers extrapolated marine practice (developed for ships) to this installation and thereby underestimated wave loading. As with the reliance on inapplicable standards, the problem solving was effectively distributed: it was not just in the designers' heads but also in the assumptions they had carried over from previous practice. Instead of examining wave loading *ab initio* they relied on previous practice to short-cut this process. Again the failure was in using something handed down in their particular culture, although this time it was not an artefact but an assumption. So this case provided an example of how an inappropriate reliance on cultural assumptions, perhaps through ignorance of their limiting conditions, can lead to failure. Cultures are obviously powerful ways of transmitting customs and practices to people in an organisation or industry, and often serve to disseminate important lessons from experience. But cultural assumptions are typically tacit, not explicit, and typically lack analytical underpinnings.

INSENSITIVITY TO ANOMALIES OF TASK DECOMPOSITION

Yet another contributor to this same capsize was that the design was modified to add a fitting to the main structure. This subsequent design step effectively violated the assumptions made in the earlier stress analysis, and the violation went unnoticed. The need to repeat the analysis also seemed to go unnoticed. Thus the design process was distributed over time, with the main structural design task being separated from a modification task, and the anomalies that arose from this were not noticed. It is perhaps quite common for failures to arise from modification, because successful modification relies on making the right inferences about the rationale of the original design - and this rationale is often obscure. But similar problems arise whenever the design task is decomposed in some way - perhaps between different individuals rather than different times. There are usually subtle inter-dependencies between decisions made in the sub-tasks, and there is no guarantee the designers involved in all the sub-tasks will be aware of these. Thus the people who design modifications may be unaware they are violating assumptions made by the designer of an original device; and the people who design, say, vessels may be unaware they are violating assumptions made by the structural designer. Obviously, wise designers will consult others in case they are violating their assumptions, but the evidence is that not all designers are wise - and in particular some are unwise about the anomalies that accompany task decomposition.

DILEMMA OF SPECIFYING BOUNDARY CONDITIONS

In one accident, in which there was a sub-sea explosion, a gas release and a subsequent fire, a hose failure was implicated. These hoses had probably exceeded their useful lives, although service lives were not specified so they had not been changed. The information processing associated with working out whether a piece of equipment is beyond its safe life is naturally distributed between designers and operators. Designers understand what is needed for the equipment to achieve functional performance. Operators have local knowledge of the service environment and pattern of use, and can often make repeated observations of the state of the equipment. This distribution, however, is naturally vulnerable to failure in certain ways. If the designer decides service life should *not* be specified to operators because the responsibility must rest with the operator to monitor an equipment's state the designer is also making assumptions about what operators can reasonably do. Either they may not easily be able to do the monitoring, or they may give it too low a priority for it to receive any resources. The designer may typically be unaware of all the other tasks the operators have to perform. On the other hand, specifying service life usually means making certain assumptions about how the equipment is treated and the environment it operates in - of which designers may be uncertain.

It may encourage operators to think they need not monitor something if all they need to do is replace it at a fixed interval, and it relieves them of making judgements about a device's condition. So, either way, whether the designer specifies service life or not, there appears to be no foolproof strategy for avoiding failure. The argument could perhaps be extended to specifying boundary conditions of any kind. Designers understand how something performs more profoundly than operators, yet the operators know the local conditions in which something is having to perform. Neither on their own can fully be aware of when something is going to fail. Yet there is often no possibility of a continual dialogue between them.

INCONSISTENCY IN INSTALLATION MODELS

In another accident, a blow-out occurred and led to a heavy spill because both a shutdown valve and a blow-out preventor were incorrectly installed. Obviously when one person or organisation designs something that another person installs there is distributed information processing and there is redundancy - two mental models of the same entity in two people's minds. The models obviously need to be consistent. There are different ways of achieving consistency, for example by writing installation manuals, but all seem to have limited power. Even when designers can physically foolproof a device there is no guarantee of consistent models. We have come across cases where people have defeated foolproofing on both electronic and hydraulic equipment by physically destroying the features that provided the foolproofing. This shows how strong mental models can be in certain cases, and how resistant they are to cues which contradict them. As with the other failure modes we have described here, there is no obvious, guaranteed remedy. But the message to, say, designers is that one of their prime objectives should be to bring the user's mental models into line with their own, and it makes sense to use a variety of means to do this since none on their own is foolproof.

INCONSISTENCY IN CHOICE MODELS

In a different accident there was a gas leak from a wellhead because of corrosion of tubing below a sub-surface safety valve. A chrome connection should have been fitted but was not. So, as in the previous paragraph, the information held by designer and installer needed to be consistent and turned out not to be. Here this replicated information was not an installation orientation but a choice of material. We do not know from the accident investigation why the installer made the wrong choice, and there could have been incentives to make the wrong choice - for example saving money, or saving time (if only non-chrome connections were available to hand). It could be the case, therefore, *not* that an installer was ignorant but that he or she gambled that failure would not occur. The motivation for the gamble was an easier task. The problem was that the gamble was ill informed - since the installer was unlikely to have had good knowledge of the probability of failure. Designers would, we think, have been able to make better probability judgements. Since operators are almost always under some kind of pressure they are almost bound to gamble. So arguably the distributed cognition problem is how to get the designers' superior knowledge of the failure odds to the operator so that the operator can make informed gambles.

DISCUSSION

GENERAL INSIGHTS

We do not feel there was anything new in our diagnoses of specific accidents, and of course we relied on other people's investigations for our data. But we do believe that the principle of distributed cognition revealed important patterns. It illustrates, for example, the importance of

'culture' - the customs, practices, assumptions and tools that provide the background to engineering - and the ways in which it can fail and then cause accidents. The development of engineering knowledge has some interesting aspects[9] which contribute both to the power of engineering and to its vulnerability. We could see, in our analyses, how the practices that led to accidents could - in only slightly different circumstances - have been seen as considerable successes. For example, carrying across a practice from one domain into another has been responsible, in part, for some major disasters like the failure of the oil platform the *Alexander Kielland*. Yet the same act of carrying across a practice from one domain to another has been seen as a considerable creative leap in some well-known product designs[6]. We need to see failure and success as radically different outcomes which can arise from virtually the same practices, we need to be attuned to exactly what makes the difference - and we need to avoid dividing up the world into people and processes that fail and people and processes that succeed.

PRACTICAL IMPLICATIONS

The development of the main practical output of this project has not yet begun, so we have not properly examined the practical implications. Moreover, looking at the failure modes we described in the Results it is hard to formulate realistic general rules or practices that would dependably avoid such accidents. Our belief is that there are no systems or tools that would always prevent these kinds of failure. But one thing that might help is for designers to have in mind these general failure modes so that they are sensitive to the possibility of them occurring. This is really the premise for the second part of our programme in which we try to develop practical tools and practical guidance on the basis of the results. These are likely to involve asking the same questions that we did in specific cases - how is information processing distributed and how does this distribution fail?

One of the difficulties with acting on our analysis, however, is that what we have identified as being the cause of failure is often also the source of productive effort. People can accomplish things in reasonable timescales because they do not have to comprehensively re-invent and re-analyse, but can draw on partial solutions given by their cultures or their environments. If you criticise someone because they do not examine the assumptions underlying their use of a standard or an existing design you have to say where the additional time needed for this examination is going to come from. Therefore the recommendation should be *not* that people avoid the sources of failure described in the Results (like 'cultural artefacts') but that they are mindful of the ways in which failure can ensue. Arguably, from the standpoint of safety, it is better to analyse exhaustively whether a particular device suits a particular application than it is to simply follow the last design that used this device. But if this is impossible then a second best situation is to follow the last design but examine 1) what conditions would make the device fail, 2) ask whether any of these conditions are more likely in the new application than the old.

LIMITATIONS OF THE WORK

There are some obvious limitations to do with the data we relied on (other people's accident reports), and because we are at an early stage of the work we have not worked out the practical implications properly. But there are also some fundamental limitations in the analysis we conducted. In particular, there is usually the possibility of explaining people's wrong decision making both from a cognitive standpoint (where they have inappropriate beliefs) and a motivational one (where they have inappropriate motivations or incentives). For example, the installer who used a plain steel fitting instead of a chrome one may have been ignorant of the corrosive fluid or may have gambled that the fitting would not fail in his

lifetime. Even when one is observing someone in the process of making an error it can be very hard to determine which of these two kinds of diagnosis is the better - the cognitive or the motivational. This means that our analyses of particular cases are always going to be debatable. But since the purpose of the work is to influence people in the future what matters is what *could* happen. And if we can argue that failures in distributed cognition could happen, and could cause accidents, our results should have some usefulness.

ACKNOWLEDGEMENTS
Many thanks are due to Bob Miles of the HSE for facilitating this work, which is being funded by the HSE under contract D3916.

REFERENCES

1. Busby J.S., 2001, Error and distributed cognition in design, *Design Studies* 22: 233-245.
2. Salomon G., 1993, Editor's introduction. In Salomon G (ed). *Distributed Cognitions: Psychological and Educational Considerations,* Cambridge University Press, Cambridge UK, xi-xxi.
3. Hutchins E., 1995, *Cognition in the Wild*, The MIT Press, Cambridge MA, p.155.
4. Norman D.A., 1993, *Things That Make Us Smart. Defending Human Attributes in the Age of the Machine,* Addison-Wesley, Reading MA, p.146.
5. Marti P., 2000, The choice of the unit of analysis for modelling real work settings, *Cognition, Technology and Work*, 2: 62-74.
6. Busby J.S., 2001, Practices in design concept selection as distributed cognition, *Cognition, Technology and Work*: forthcoming.
7. Lave J., 1988, *Cognition in Practice*, Cambridge University Press, Cambridge UK.
8. Scaife M. and Rogers Y., 1996, External cognition: how do graphical representations work? *International Journal of Human-Computer Studies,* 45: 185-213.
9. Blockley D.I. and Henderson J.R., 1980, Structural failures and the growth of engineering knowledge, *Proc. Institution Civil Engineers Part 1*, 68: 719-728.

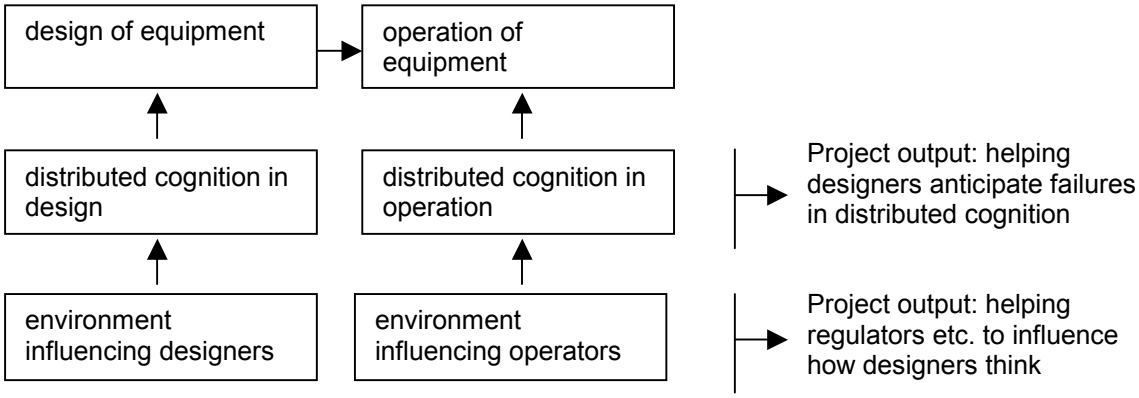

Figure 1: Scope of the project