

INTERFACE CHANGES GENERATING ACCIDENTS. A SCHEMA-DRIVEN APPROACH OF NEGATIVE TRANSFER.

CACITTI, Lucile
Electricité de France
Département Facteurs Humains
41220 Saint-Laurent des Eaux
FRANCE
&
BESNARD, Denis
Department of Computing Science
University of Newcastle
Newcastle upon Tyne, NE1 7RU
UNITED KINGDOM.

Abstract. When expert operators interact with a new device, they inevitably reuse former interaction modes and actions. This phenomenon is due to the human cognition seeking resources savings. Schemas support this strategy and are implemented in such a way that perfection is disregarded at the profit of an intuitive trade-off between performance and cognitive resources savings. As a consequence, humans have a strong inclination to fit well-known solution procedures into new problems. For this reason, changes in work environments can cause accidents when they allow operators to erroneously interact with a new device if the latter is perceived as familiar. This research issue originates from an industrial background. The suspected cause of a fatal error performed by an operator in a steelworks factory is replicated in a simple experiment. The results support the hypothesis according to which errors (and eventually accidents) due to changes are more likely when the latter do not inhibit irrelevant former interaction modes. This main result is discussed in the light of cognitive psychology.

1 INTRODUCTION

The study of field situations from the standpoint of cognitive ergonomics aims at understanding cognitive acts within the context in which they happen. Humans, their tools, their reasoning processes and actions inside the environment are classical features of this kind of approach. The latter can be deliberately quantitative when research aims at isolating a particular parameter, e.g. the cause of an error. In this case, experimentation can be used to assess the effect of one or several factors on a given aspect of behaviour. This is the option taken in this paper. A field study has been conducted in a steelworks company where an accident occurred, leading to the death of an operator. This study has been initiated in order to trace back the psychological causes of this accident. The latter will be treated as errors in the human-machine dialogue during exception handling.

We identify two wide classes of exception in human-machine interaction: a) exceptions that occur after deployment for which designers have not conceived any procedure due to the unlikelihood of these events and b) situations that are unexpectedly similar to others for which a well-defined procedure exists. Clearly, our paper deals with the second case and will try to highlight the risks associated with certain types of similarities at the interface level.

Our approach relies on cognitive ergonomics for at least two reasons. Firstly, the knowledge of the causes of the accident would be incomplete if the mental processes of the operator could not be assessed. Safety and ergonomics at the workplace have to take into account the psychological aspects of the tasks that humans have to perform (Christol & Mazeau, 1991). Secondly, cognitive ergonomics provide a theoretical and methodological framework sitting on several decades of cognitive psychology. This allows cognitive ergonomists to carry out research about mental processes under experimental settings, if needed. This is the case in this research.

The present paper will rely on a psychological theoretical framework in order to document the core factors involved in the accident. We will then test these factors in an experiment where we wish to assess the mental processes involved. For scope matters, we will clearly disregard organisational factors, although we acknowledge they always play a significant role in accidents (see Reason, 1990; 1995; 1997; 2000; Bieder, 2000). In doing so, we will miss the fruitful interaction of a multi-layered analysis. On the other hand, it will allow us to allocate more effort in an in-depth study of individual factors.

The outline of the paper is as follows. We will start by presenting cognitive concepts for framing the research (section 2). The article will then briefly outline the accident (see section 3) and the method of the experiment (section 4) designed for testing its suspected causes. The results (section 5) will lead to a discussion on the theoretical and practical outcomes of our research (section 6).

2 SCHEMA-DRIVEN HUMAN COGNITION

2.1 The concept of schema: an overview

Schemas have been studied or mentioned in a wide class of papers whose topics include medical diagnosis (Lesgold *et al.*, 1988), car driving (Van Elslande, 1992), problem solving by analogy (Catrambone & Holyoack, 1989; Novick & Holyoack, 1991), aircraft piloting (Amalberti, 1992) and computer program understanding (Detienne, 1996). Schemas are high-level knowledge structures that support any aspect of knowledge and human skills (Reason, 1990). They support the fast processing of routine situations for which one acts virtually automatically from their identification. The concept of schema is close to Rasmussen's (1986) rule-based level of control and the parallel has already been established (Bollon & Channouf, 1993)¹. For the time being, let us just assert that a schema-based action is conditioned by the identification of a set of activators in a situation (e.g. the statement of a problem or the symptoms of an illness). These activators then trigger the schema which, in turn, controls the actions performed. The process is roughly similar to an IF...THEN statement where some conditions have to be detected for the schema to trigger (Govindaraj & Su, 1988). Historically, the concept of schema originates from Bartlett (1932) but some psychological processes similar to schemas have been described, among others, by such terms as experiential knowledge (Fink & Lusth, 1987), scripts (Shank & Abelson, 1977; Boshuizen *et al.*, 1991; Custers *et al.*, 1996) and frames (Minsky, 1986). As we have stated, a schema is a piece of knowledge meant to solve familiar problems. But as all problems are not familiar in the first place, we first need to have a look at a potential explanation for schema building. This will feed our description of the nature and role of schemas. When a problem is unknown², one tries to solve it by trial and error, or formally speaking, by hypothesis testing (Byrne, 1989; Liu, 1991). Once a solution has been found, it usually can be stored in memory and thus becomes repeatable. With time, the repetitive exposure to the same kind of problems leads to the recall of a generic solution and the progressive building of triggering rules that bind together a) the solution and b) the type of problems it solves. On the basis of experience, this set of rules is progressively refined and tuned until it triggers the schema only for the relevant cases (see Figure 1 below). But because humans found their interaction with the world on memory of past experiences (Randel & Pugh, 1996; Roediger, 1980), it happens that for most of new problems, there is a solution to another problem that can be adapted for reuse. This is a very common process in human cognition but as we will see in the next section, it is also a fallible one.

¹ Although schemas could be matched with the skill-based level of Rasmussen's (1986) model, we object that the latter is more about a sensorimotor level of control. We thus prefer analogy to the rule-based level of control, following Salminen and Tallberg (1996).

² e.g. the Hanoi tower, assuming one has never solved anything similar before.




Figure 1: Knowledge refinement can be compared to a ball rolling in a more and more narrow ditch (adapted from Rauterberg, 1995)

2.2 Schemas and transfer.

So far, schemas have been presented as a form of knowledge that supports human reasoning. As this research is interested in the cognitive features involved in the occurrence of an accident, the fallible aspect of human reasoning will be addressed here. We will thus revisit schemas under a very common angle in cognitive ergonomics: error. But in doing so, we have to emphasize that human errors are not always mere cognitive dysfunctions. Often, and it is the case in our study, errors are marginal events caused by the same mechanisms that generate correct acts most of the time (Johnson *et al.*, 1992). As a consequence, errors are not by-products of cognition. They are the side-effects of a risk induced by a heuristic reasoning process, the latter being aimed, time after time, at getting an optimal performance for the lowest mental cost (Amalberti, 1996).

Since Simon (1957) and his concept of bounded rationality, it is accepted that humans' actions do not reach perfection. In this conception, humans' actions rather seek optimality with respect to their goals and what the cognitive resources allow. The fact that the cognitive system is not aimed at handling all the data available in the environment is a central aspect of the cognitive resources saving strategy. Humans select the data that, by experience, have been discovered to be required for a given task. Furthermore, these data do not all have the same saliency. The ones that support a core function of a task are often those which must be detected in priority. They happen to be stored in memory with functional alterations in order to reflect their specific role (Ochanine, 1978; Moray, 1987; Endsley & Smith, 1996). This view is part of a modern operational formalisation of reasoning activities where cognitive resources support the fallible execution of a task rather than the exhaustiveness of a pure logical analysis.

In line with the resources saving strategy, schemas provide ready-made solution procedures in response to a situation identified as a pattern of data rather than as a series of hierarchical goals (see Boreham, Foster & Mawer, 1992), the latter requiring a higher level of control and higher resources involved. Among drawbacks with schemas (see Reason, 1987a for an overview) is that they can be triggered as soon as a known pattern of data is detected, even if the latter is incomplete. Thus schemas can lead to errors when such a pattern is detected in an unknown problem. This problem can then be recognised as familiar, processed like a routine one with high probability of errors. This explains why practitioners sometimes have difficulties in identifying exceptional diseases and confuse their symptoms with more benign cases. Moreover, as expert operators usually allocate few resources to the execution of a schema-driven action, they have difficulties in detecting exceptions. This phenomenon, experimentally studied by Besnard (1999; 2000) and Besnard & Cacitti (2001) is a potential explanation for errors committed by expert operators.

The heuristic cognitive acts allowed by schemas aim at saving the resources allocated to the execution of a given task. One of the weaknesses of this saving strategy lies in the potentially flawed management of changing situations, which is the scope of this paper. When a given situation is such as its changes can be overlooked, it allows one, especially an expert operator, to trigger an irrelevant schema. Then a negative transfer can occur that alters the execution of the task. It can be a potential cause of accidents when it occurs in an environment where safety is a critical dimension. To us, several conditions are needed for this negative transfer to occur:

- The operator must be experienced since the transfer we report on in this research is caused by a domain-specific (Schanteau, 1992) schema or a rule-based level of control;

- The change in the situation must be such that the latter can be treated as another well-known situation.
- The usual underlying structure of the current type of problem or situation must be discrepant to the surface features used to support the human mental model.

As the origin of the study lies in the analysis of an accident that occurred in a steelworks factory, this research is embedded in an ergonomic framework that aims at analysing the workplace using concepts about operators' mental processes. Thus, after having reviewed some of these concepts, we are now going to the core of the paper and consider cognition in an industrial context where the details of the accident will be exposed.

3 DESCRIPTION OF THE ACCIDENT

The following accident occurred in 1991 in the south of France, in a major steelworks factory employing some 500 people. An experienced operator was working on a thread drawing machine, a device that reduces the diameter of a metallic thread by a series of tractions. The output thread is coiled onto a drum and kept in place by pressing pads. On this specific machine, the lever position for opening and closing the pads was swapped as compared to the eight other machines. This swap was well-known by the operators but was not flagged or equipped with any kind of protection. Because of the swapped commands, the operator has opened the pads whereas his intention was to close them. This error occurred at a time of the process where opening the pads is forbidden. The operator was violently hit by the thread uncoiling itself from the drum. This accident has been fatal.

The tools' characteristics were discrepant with respect to the routine control mode. The operator was accomplishing a routine task with a tool he did not recognise as an unusual one. This contributed to the occurrence of the fatality. In the accident exposed here, the schema implemented by the operator did not match the usual constraints imposed by the tool. As a consequence, a routine schema controlled a task in which an exception had not been detected. Actually, the accident did not only occur because the operator performed an error (see Doireau, Wioland & Amalberti, 1995) but because the conditions in which this error occurred were unusual. When a tool changes, e.g. for an upgrade, the schema must change accordingly in order to reflect the changes and maintain the accuracy of the interaction. But updating such a complex knowledge structure is not performed by merely acknowledging that the tool has changed. Instead, it requires repetitive feedbacks from the system in a wide variety of cases. Progressively, the discrepancies between the expected system's behaviour and the actual system's behaviour drive the updating process of the schema. So on the one hand, the operator's errors progressively contribute to refine the performance. But on the other hand, the operator's knowledge will stay partly inaccurate as long as all the discrepancies have not been detected and accommodated for. During this sensitive period, errors on critical functions of a hazardous tool can be fatal.

Our objective is now to investigate some of the industrial accident factors in laboratory settings. As Green and Hoc (1991) and Hoc (1993) suggest, this is a classic and fruitful approach in cognitive ergonomics. Although it could be objected that lab experiments are far too reductive as compared to the complexity of natural environments (as suggested by Perruchet, 1997), it nonetheless originates from a field situation, giving some credit to our approach (Sperandio, 1995) and some soundness to our desire to carry out this research. Among others, its industrial background answers the question of "why" we want to conduct this work. Lab experiments also allow one to isolate a specific factor and to study it without unwanted contextual side-effects (work colleagues' conversation, unavailability of operators as experimental subjects, managers supervising the operator during the experiment, etc.). Moreover, simulation permits to replicate errors with little concern about their immediate consequences for the experimental subject. Lastly, and this comment is a rather epistemological one, even small-scale experiments are worth attempting. Since psychological data, however microscopic they are, can be interpreted in a usually well-documented theoretical background, they still permit to increase the predictive power of psychology. This approach is one where the human cognitive system is considered as a deterministic machine but whose complexity is still beyond our current predictive

capacities. That is not the case for disciplines such as mechanics or electrical engineering (according to Life *et al.*, 1996) or even computing science. But it does not mean that such a state of knowledge (a high predictive power) cannot be reached. It may only be a matter of time, assuming that it is the objective we want to reach and that we have the right methods.

In line with the accident described above, and the theoretical concepts exposed in section 2, we expect errors to occur as a function of the similarity of a new interface to the former one. We know from the literature that if a schema can detect, in a changed device, some of the activators it needed to trigger for a former device, it is then likely that this schema will then control the interaction with the new device. So if some discrepancies exist between what the behaviour of a new device is and what the schema prescribes as correct actions, then errors (which potentially would have been correct actions on the former device) have to be expected. This hypothesis will now be tested on a simple simulated control task.

4 METHOD

4.1 Description of the task

Because we focus on negative schema transfer under experimentally controlled conditions, we have defined a simple computer-based control task (see screenshot Figure 2). In this task, the subject has to fill up 4 classes of containers (upper left corner) with four classes of items (middle boxes).




Figure 2: Screenshot of the interface of the control task

We have chosen to use fruits (bananas, pears, cherries and strawberries) for their ease of identification and intuitive meaning. The central area of the screenshot displays the four areas where the four classes of items are represented. The bottom count bar displays the number of items left to be treated.

The task is cyclic and is composed of three stages.

- *Step 1, Pointing.* The subject asks the system to randomly point to a class of items. A black bar is then displayed on the top of the class box which the system points to.
- *Step 2, Selecting.* The subject has to respond by selecting the same class of items as the one pointed by the system. The class selected by the subject is highlighted in the bottom count bar.

- *Step 3, Filling or Emptying.* The subject asks for one item of the selected class to be sent in the corresponding box. At this point, the cycle goes back to stage 1. When a box contains 3 items, the subject must ask for it be emptied and a container appears on the screen. The cycle then goes back to step 1 (see Figure 3 below).

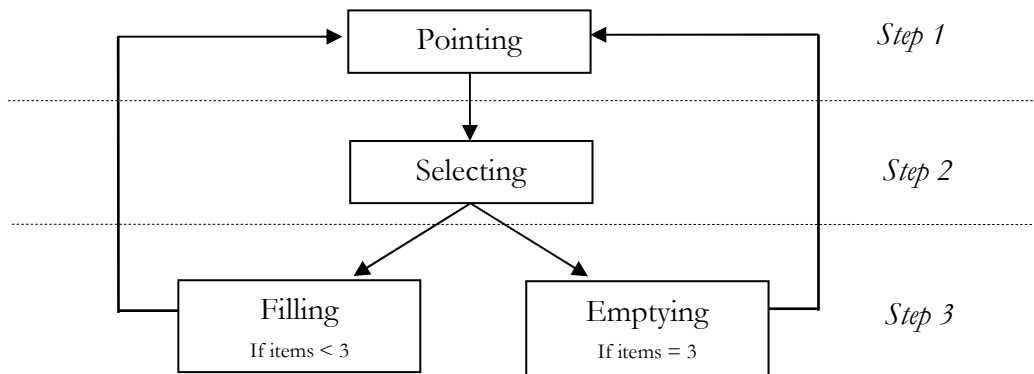


Figure 3: Description of a cycle


4.2 Subjects

Twenty French students, unselected for age and sex, from various departments of the University of Provence joined the experiment. They all had the same minimum skills on the system due to a training phase where they all complied to a performance criteria.

4.3 Training phase

During this phase, the subjects had to execute 108 cycles³ without error with the interface displayed in Figure 2. The task was restarted if an error occurred. The controls were keyboard-based according to Table 1. The keys had colour stickers on so that subjects could easily locate them on the keyboard.

Table 1: Key-function mapping for the training phase

Key				
Function	Pointing	Selecting	Filling	Emptying

4.4 Experimental conditions



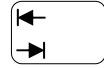
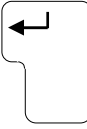
After the training session, the subjects were assigned to one of the two following experimental conditions where they had another 108 cycles (still in 3 trials) to perform:

- *Swapped* commands. The controls were keyboard-based according to Table 2.
- *On-screen* commands. The controls were icons displayed on the screen (see Figure 4 for a screenshot) and mapped to control functions as shown in Table 3.

In each condition, an on-screen message was displayed when a subject made an error. This had no consequence over the performing of the task.

³ The decomposition is as follows: 9 items per class x 4 classes x 3 trials.

Table 2: Key-function mapping for the *swapped* condition

Key				
Function	Pointing	Selecting	Filling	Emptying

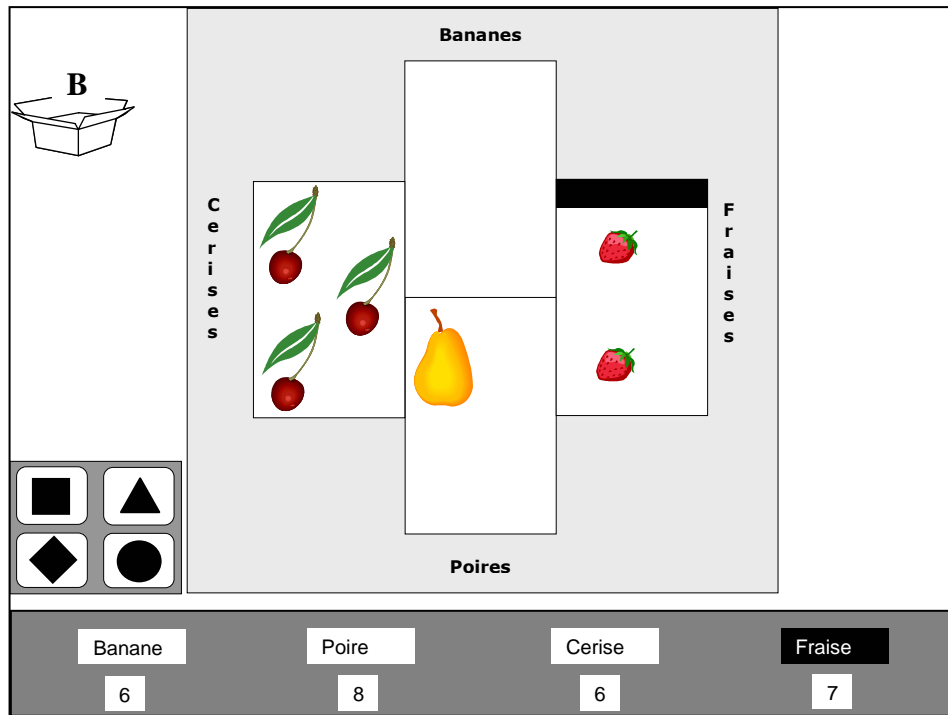






Figure 4: Screenshot of the interface for the *on-screen* condition

Table 3: Icon-function mapping for the *on-screen* condition

Icon				
Function	Pointing	Selecting	Filling	Emptying

4.5 Variables

The subjects performed 3 trials from which means were computed and formed the following dependent variables⁴:

- *Time*. This is the number of seconds needed by the subjects to complete the task.
- *Total*. This is the mean of all the other variables below. Because it is a mean, its value is lower than some of the variables.
- *Omission errors*, where one or several steps are skipped in a cycle.
- *Commission errors*. It is an action that is not relevant to the current system's state. It is the case of a subject who would empty a class box whereas the latter is not full (i.e. it contains less than 3 items).
- *Previous interface*. These errors would have been correct actions under the training interface. This is the explicit variable for our analysis of the negative transfer.

⁴ The variables labelled *omission errors* and *commission errors* are derived from Gobet & Simon (1996) and Hollnagel (1993).

- *Other errors.* These are erroneous actions such as mistyping or any other action that cannot be interpreted as belonging to the above variables.

4.6 Predictions

We assume that errors will originate in the failure to inhibit key-function couplings built during the training phase. When the interface changes, we expect interferences to occur due to the persistence of these previous couplings. As our aim is to experimentally investigate the negative transfer rather than decomposing into its sub-components, we do not make predictions for each variable. We will only expect the following:

In the *swapped* condition, we globally expect a large number of errors since the same keyboard keys are now dedicated to different functions. The similarity with the key-function mapping of the training phase should leave enough room for the former schema to partly override the learning of the new interface. This may cause major disruptions in subjects' performance and it is the condition where errors due to the previous interface (the training one) are expected to be highest.

In the *on-screen* condition, the very nature of the interface has changed. The subjects now control the system via a mouse by clicking on icons displayed on the screen (see screenshot in Figure 4). Because of the difference between the training interface and the current one, we predict that this condition restricts the possibilities of transferring previous key-function couplings. As a consequence, we expect less disruptions in this condition. Additionally, we expect the errors due to the previous interface to be lowest in this condition.

5 RESULTS

The significant results are summarised in Figure 5 and in Table 4. The comments will appear in the next section.




Figure 5: Graphical representation of the significant means

Table 4: Summary of the significant means

	total	omission	commission	prev. interface	other
on-screen	0.3	0.03	0.23	0.6	0.3
<i>SD</i>	<i>0.33</i>	<i>0.10</i>	<i>0.22</i>	<i>0.7</i>	<i>0.7</i>
swapped	1.8	1.2	0.57	2.7	2.7
<i>SD</i>	<i>1.29</i>	<i>1.98</i>	<i>0.47</i>	<i>1.9</i>	<i>2.7</i>
<i>F value</i>	<i>12.54</i>	<i>9.4</i>	<i>4.05</i>	<i>10.87</i>	<i>7.6</i>
<i>df</i>	<i>18</i>	<i>18</i>	<i>18</i>	<i>18</i>	<i>18</i>
<i>p</i>	<i>0.0023</i>	<i>0.0067</i>	<i>0.0593</i>	<i>0.004</i>	<i>0.013</i>

The effect of interface on time did not reach significance. (*on screen*=220.26; *swapped*=216.66; $F(1;18)=0.02$; $p=.885$).

The effect of interface on total errors reached significance: In the *on-screen* condition, subjects performed a mean of 0.3 errors *vs.* 1.8 in the *swapped* condition ($F(1; 18)= 12.54$; $p=.002$). Three out of five of the error measures (omission, commission & previous interface, respectively) reached significance. The details of these three results now follow. Firstly, the number of omission errors differed significantly across groups: In the *on-screen* condition, subjects performed 0.03 errors *vs.* 1.2 in the *swapped* condition ($F(1;18)=9.40$; $p=.006$). Secondly, the effect of interface on commission errors produced significant differences: The *on-screen* subjects performed 0.23 error *vs.* 0.56 for the *swapped* subjects ($F(1;18)=4.05$; $p=.059$). Thirdly, as predicted, there were fewer transfer errors in the *on-screen* condition ($m=0.6$) than in the *swapped* condition ($m=2.7$) ($F(1;18)=10.87$; $p=.004$). Similarly to the above, the effects of the interface change on the other errors were lower in the *on-screen* condition ($m=0.3$) than in the *swapped* condition ($m=2.7$; $F(1;18)=7.60$; $p=.013$).

6 DISCUSSION

As we have seen in the previous section, all the significant error measures show higher values in the *swapped* condition. This is the interface where negative transfer effect is strongest and causes major disruptions. Entering in some more details, the results show that omission errors are more contrasted across the two conditions than commission errors, the latter reaching a poor level of statistical significance ($p=.059$). The errors due to the previous interface show a large significant difference between the two conditions. Again, the *swapped* condition shows the highest number of errors. The variable called *other errors* also shows higher error rates in the *swapped* condition, conforming to the trend of results. In a future work, this variable would be worth splitting and investigating more deeply as it may reveal some nuances that our high-level variables have probably masked. In our interpretation, the results are due to the fact that the activation of a former irrelevant schema is left possible in this condition. Thus, as a preliminary conclusion, it seems highly plausible that a negative transfer occurs when two interfaces or problems share surface features but have different structures (Novick, 1988; Blessing & Ross, 1996).

The following discussion will now try to bounce from the experiment's results back into "real life" considerations. In doing so, we will adopt Suchman's (1987) views on situated action according to whom a) cognitive phenomena are related to artifacts and actions b) and the significance of these artifacts and actions is related to the circumstances in which they occur.

Human cognition is extremely good at modelling the regularities of the past, storing and reusing them as a basis for some automatic control of actions (Reason, 1987b). Providing some details on this phenomenon, Reason (1990) and Decortis (1993) suggest that humans globally obey two heuristics: frequency and similarity. The frequency heuristic assumes that an operator processes any situation by applying the solution that has, in the past, generated correct actions most of the time. Let us take the example of a doctor performing a diagnosis: He or she will interpret the patient's symptoms according to the explanation that has historically revealed correct most often (Patrick, 1993). This behaviour is called heuristic because it relies on intuitive probabilities and induces an implicit risk. At least two of our previous studies on expert trouble-shooters' errors (Besnard & Bastien, 1999; Besnard, 2000) have

given some credit to Reason's frequency heuristic. Our research now experimentally highlights the reality of the similarity heuristic through another example of schema-based error. People, when they perceive some familiarity in a novel problem (or when they do not detect an exception), tend to solve it by applying a pre-existing solution belonging to a category of problems that seems to be similar to the current one. Again, this behaviour is heuristic since it integrates an implicit risk: Similarity is based on surface features that can be detected incompletely without this incompleteness being noticed. To this respect, the cognitive resources saving strategy pushes for a similarity-based matching and can cause exceptions to be overlooked.

In hindsight, the changes we have implemented in our experiment can be seen as dramatic ones. If one considers the dialogue point of view, the *swap* condition has a huge impact on the interaction. They allocate previous commands to new functions and as demonstrated by the experiment, this is the worst change one could ever think of. But to some extent, these changes are not that radical from a design point of view. The commands and the functions in the *swap* condition are still the same as in the training phase, as opposed to the *on-screen* condition. So the issue here is that under changing conditions, opting for a totally different design can increase the level of dependability of a given interaction as opposed to modifying an interface. Now, adopting a contextual point of view allows us to reach another dimension. Suchman's (1987) views suggest that the context has to be taken into account for qualifying the changes. And the context is that the commands swap was implemented on only one machine out of eight. So it may be the discrepancy of the interfaces across the machines, more than the mere isolated commands swap, that caused the accident. In this conception we think that interface consistency is a serious issue when an operator interacts with several machines.

As we have seen before, it is precisely at the surface level that the schema-driven, heuristic pattern-matching behaviour takes place. Since this process can occur with little regard to the underlying information or system structure, surface changes potentially have huge effects. Even researches claiming that experts focus on the structure of data for solving problems (see Hardiman, Dufresne & Mestre, 1989; Smith, 1992; Zajkowski & Martin, 1993) do not demonstrate that data are detected other than as a pattern. As an information source, the surface is preferred (Blessing & Ross, 1996). Consequently, the more a new interface displays surface features that are similar to a well-known interface, the more likely it is that the same schema will be used for guiding the interaction. This is why we think having a variety of interfaces on similar machines is potentially dangerous in critical environments. A possible explanation lies in the fact that human cognition forces human-machine interaction into a mode where surface similarity and reuse of knowledge are strong dialogue drivers.

It now has to be said that the heuristic interaction modes that humans adopt in their routine actions usually provide an acceptable level of reliability. This strategy only fails when exceptions are not detected. Then, the discrepancy between what the situation is and what the operator thinks it is, is where risk lies. So we must make clear that the similarity heuristic is not an error generator *per se*. It only is so when it triggers actions whose discrepancy from the optimal interaction mode is overlooked.

In our opinion, our experimental settings are representative of the aspect we wanted to study from the accident: Knowledge is transferred across tools and interfaces and this sometimes happens out of any control from the subject. It can nonetheless be objected that our experiment involves a tool and subjects that are very different from the original industrial situation and this issue is a serious one (Karnas & Van De Leemput, 1990). However undeniable this is, we do think our study is focussed on a central psychological mechanism that we think is not situation specific.

6.1 Errors caused by expertise

Expert operators can cross the boundaries of their expertise without awareness. In the accident described in section 3, the fatality was caused by the operator behaving with a new tool the way he used to behave with the other ones. The fact that the properties of the tool had changed without these changes being accounted for is the issue. The operator had not identified the exception at his workplace and implemented routine actions under misdetected non-standard settings. We will now be discussing this phenomenon under a rigidity angle.

It is already known that the more expert an operator is, the more rigid (i.e. less adaptable) his or her knowledge tends to be (Rasmussen & Jensen, 1974; Moray, 1987; Hollnagel, 1987; Gaba, 1991). As Reason (1987a) puts it, schema-based reasoning is rigid and rule-bound: Solutions to previous problems can thus be applied with little attention paid to changes. Rigidity is a drawback especially in new situations. It leads to poor adaptability because of the schema-driven or rule-based interaction mode being highly prevalent upon any other. Operators then ignore information which does not fit in the active schema (Air Inter, 1995) and/or repeatedly perform the same actions without any gain in understanding of the situation (Bereiter and Miller, 1989). These two erroneous behaviours partly compose what has been called fixation errors or cognitive lockup. The interesting aspect of these schema-based errors is that they concern high degrees of expertise. In this respect, the data gathered about the accident and the results obtained in the experiment are other pieces of evidence of the fallibility of expert reasoning. They also provide more empirical data on a mechanism that is well-known in the ergonomic literature but rarely quantitatively investigated through experiments.

Many other examples of erroneous knowledge transfer could be found. A simple one has to do with computer operating systems. By default, when Windows™ users copy a file to a floppy drive, they usually expect the files to be copied immediately. This is not true for the Mandrake distribution of Linux, and potentially, for other distributions either. Under Mandrake at least, and this is a default setting, floppy drives have to be mounted so that they can be accessed. Before they are removed, these medium have to be unmounted so that all the operations that have been performed on these drives can actually occur. If a floppy is ejected from its drive without being unmounted, the user's operations may still be cached and there is a high risk of data loss. From the results of our experiment, we can predict that - and explain why - a user who would have learned computing under the Linux Mandrake operating system would make less of such errors than another user who would have to transfer his skills from another platform. This prediction would have to be confirmed by experimentation but the leitmotiv here is to say that expertise is an imperfect state of knowledge that only reaches maximum performance in standard situations, should the latter be very complex.

Generally speaking, errors are progressively reduced by learning. By gathering repetitive feedbacks from the interaction with a system, individuals refine their knowledge and increase their performance. Through experience, the conditions needed for some specific action to trigger are progressively narrowed down to a very accurate set. This is how changes are accommodated for by humans. Having said that, we must not assume that humans, especially operators in hazardous processes, always have an opportunity for learning: they can die from their errors.

6.2 Implications at the workplace.

The accidental situation we have analysed in this research was not different enough from routine ones for them to be interpreted as exceptional. What could have caused them to be treated as such is worth spending some effort. This is the topic of the present section where we will more generally deal with counter-measures. Generally speaking, we think our results have to do with the way man-machine interaction is designed. The following suggestions derive from this assumption.

- *Knowledge evolution takes time.* When a tool or interface changes, e.g. for an upgrade, the operator has to modify his or her knowledge in order to adapt it. This requires repetitive feedbacks from the new device until some acceptable level of performance is reached. But as Aberg and Rimmö (1998) suggest, procedural knowledge, as opposed to declarative knowledge, takes time to be improved. Just as a schema and its set of rules are progressively refined so that they are used only to generate correct actions, modifying them costs time and is error-prone.
- *Compensating takes more than simply acknowledging.* When a change has to be accommodated for by operators, formal education or warning instructions are a starting point but help very little since declarative knowledge, to say the least, is not strongly correlated to the level of performance (Schraagen & Schaafstal, 1996). Our experiment demonstrated, among others, that knowing that a change has been performed is not enough for it to be compensated for.
- *Don't change the interface.* Similarity can cause accidents. This is a serious issue in critical environments. So we suggest that where the content of the dialogue is kept unchanged during e.g. an upgrade, the interfaces should be kept unchanged or at least consistent across platforms.

If the interface has to be changed or cannot be made consistent with the existing ones, one should consider the next bullet point.

- *Change things a lot rather than a little.* Under the conditions that we have studied, changing an interface can impact on the interaction dependability more than a totally different design option. So whenever there is a doubt on whether interfaces similarity is desirable, things should be kept significantly different.
- *Make changes salient.* Critical changes have to be made salient, by implementing enabling actions for instance. Saliency is an environmental feature that humans process well. The idea here is to make a break point intrude inside a routine procedure. Some advantage can be gained from it as it introduces a higher-level mode of control allowing the revision of the mental model the operator is using.

7 LIMITS

There is a number of issues that have not been addressed in this paper. One is the mode confusion angle that John Rushby and collaborators (Crow *et al.*, 2000), among others, has documented. This angle allows one to understand how correct actions in particular settings happen to be incorrect in others. We did not mention either such factors as lack of attention or slips which are known to contribute significantly to accidents. Last but not least, some management considerations could help understanding the mechanism that led to neglect interface issues in such a hazardous environment as steelworks.

8 CONCLUSION

In this paper, we attempted to understand and assess in an experiment the psychological causes of an accident that occurred in a steelworks factory, causing the death of an operator. Among other factors that we briefly list in section 7, this fatality was caused by some changes at the workplace not being taken into account by the operator during routine actions. In our opinion, a negative transfer caused these familiar operations to trigger within work settings where they no longer were relevant. The experiment supported this hypothesis: Changes in the interface of a simulated control task generated the negative transfer, causing errors due to well-known actions being called in new settings. From the accident data, the results of the experiments and the theoretical background, we concluded that human cognition fallibility accounts for some of the errors performed during changes in work settings. However, this does not mean that this state of facts has to be passively accepted. Instead, we formulate simple yet design-centred comments considering human cognition with regards to hazardous systems conception. Lastly, following Hollnagel (1993), we think the systematic study of erroneous actions has the potential to provide a better knowledge of human failure modes and to influence the design of more reliable systems. This paper represents our modest contribution to this research avenue.

9 REFERENCES

- ABERG, L. & RIMMÖ P-A (1998). Dimensions of aberrant driver behaviour. *Ergonomics*, 41, 39-56.
- AIR INTER (1995b). Détecter les situations dangereuses. *Bulletin de la Sécurité des Vols d'Air Inter*, 46, 8-13.
- AMALBERTI, R. (1992). Safety and process control : An operator centered point of view. *Reliability Engineering and System Safety*, 38, 99-108.
- AMALBERTI, R. (1996). *La conduite de systèmes à risques*. Paris, Preses Universitaires de France.
- BARTLETT, F C. (1932). *Remembering: An Experimental and Social Study*. Cambridge, Cambridge University Press.
- BEREITER, S. R. & MILLER, S. M. (1989). A field-based study of troubleshooting in computer-controlled manufacturing system. *IEEE Transactions on Systems, Man and Cybernetics*, 19, 205-219.

- BESNARD, D. & BASTIEN-TONIAZZO, M. (1999). Expert error in trouble-shooting. An exploratory study in electronics. *International Journal of Human-Computer Studies*, 50, 391-405.
- BESNARD, D. (2000). Expert error. The case of trouble-shooting in electronics. Proceedings of the 19th international conference *SafeComp2000*, Rotterdam, Netherlands (pp. 74-85).
- BESNARD, D. & CACITTI, L. (2001). Troubleshooting in mechanics: A heuristic matching process. *Cognition, Technology and Work*, 3, 150-160.
- BIEDER, C. (2000). Comments on the JCO accident. *Cognition, Technology & Work*, 2, 204-205.
- BLESSING, S. B. & ROSS, B. H. (1996). Content effects in categorization and problem solving. *Journal of Experimental Psychology : Learning, Memory and Cognition*, 22, 792-810.
- BOLLON, T. & CHANNOUF, A. (1993). Travail sur machine-outil et activité de planification. Actes du colloque de prospective : *Recherches pour l'Ergonomie*. Université de Toulouse-le-Mirail.
- BOREHAM, N. C., FOSTER, R. W. & MAWER, G. E. (1992). Strategies and knowledge in the control of a chronic illness. *Le Travail Humain*, 55, 15-34.
- BOSHUIZEN, H. P. A., HOBUS, P. P. M., CUSTERS, E. J. F. M. & SCHMIDT, H. G. (1991). Cognitive effects of practical experience. in D. A. EVANS & V. L. PATEL (Eds) *Advanced models of cognition for medical training and practice*. Heidelberg : Springer Verlag. (pp. 337-348)
- BYRNE, R. M. J. (1989). Human deductive reasoning. *The Irish Journal of Psychology*, 10, 216-231.
- CATRAMBONE, R. & HOLYOAK, K. J. (1989). Overcoming contextual limitations on problem-solving transfer. *Journal of Experimental Psychology : Learning, Memory and Cognition*, 15, 1147-1156.
- CHRISTOL, J. & MAZEAU, M. (1991). Ergonomie cognitive : pléonasme ou paradoxe ? *Le Travail Humain*, 54, 381-390.
- CROW, J., JAVAUX, D. & RUSHBY, J. (2000). Models and mechanized methods that integrate human factors into automated design. In proceedings of the *International Conference on human-computer interaction in aeronautics : HCI-Aero 2000*, Toulouse, France, September 2000.
- CUSTERS, E. J. F., BOSHUIZEN, H. P. A. & SCHMIDT, H. G. (1996). The influence of medical expertise, case typicality and illness script component on case processing and disease probability estimates. *Memory and cognition*, 24, 384-399.
- DECORTIS, F. (1993). Operator strategies in a dynamic environment in relation to an operator model. *Ergonomics*, 36, 1291-1304.
- DETIENNE, F. (1996). What models for program understanding? Actes du colloque *Using Complex Information*. Université de Poitiers (pp. 60-66).
- DOIREAU, P., WIOLAND, L. & AMALBERTI, R. (1995). La détection de l'erreur par un tiers en situation de pilotage d'avions. *Service de Santé des Armées. Travaux Scientifiques*, 16, 291-292.
- ENDSLEY, M. & SMITH, R. P. (1996). Attention and decision making in tactical air combat. *Human Factors*, 38, 232-249.
- FINK, P. K. & LUSTH, J. C. (1987). Expert systems and diagnostic expertise in the mechanical and electrical domains. *IEEE Transactions on Systems, Man and Cybernetics*, 17, 340-349.
- GABA, D. (1991). Dynamic decision making in anesthesiology : Cognitive models and training approaches. in D. A. EVANS & V. L. PATEL (Eds) *Advanced models of cognition for medical training and practice*. Heidelberg : Springer Verlag (pp 123-147).
- GOBET, F. & SIMON, H. A. (1996). Templates in chess memory : a mechanism for recalling several boards. *Cognitive Psychology*, 31, 1-40.
- GOVINDARAJ, T. & SU, Y. D. (1988). A model of fault diagnosis performance of expert marine engineers. *International Journal of Man-Machine Studies*, 29, 1-20.

- GREEN, T. R. G. & HOC, J. M. (1991). What is cognitive Ergonomics ? *Le Travail Humain*, 54, 291-304.
- HARDIMAN, P. T., DUFRESNE, R. & MESTRE, J. (1989). The relation between problem categorization and problem solving among experts. *Memory and Cognition*, 17, 627-368.
- HOC, J.-M. (1993). Conditions méthodologiques d'une recherche fondamentale en psychologie ergonomique et validité des résultats. . *Le Travail Humain*, 56, 171-184.
- HOLLNAGEL, E. (1987). Information and reasoning in intelligent decision support systems. *International Journal of Man-Machine Studies*, 27, 665-678.
- HOLLNAGEL, E. (1993). The phenotype of erroneous actions. *International Journal of Man-Machine Studies*, 39, 1-32.
- JOHNSON, P. E., GRAZIOLI, S., JAMAL, K. & ZUALKERNAN, I. A. (1992). Success and failure in expert reasoning. *Organisational Behavior and Human Decision Processes*, 53, 173-203.
- KARNAS, G. & VAN DE LEEMPUT C. (1990). L'ergonomie de conception du dialogue homme-ordinateur. Généralité et relativité des résultats. *Le Travail Humain*, 53, 303-311.
- LESGOLD, A., GLASER, R., RUBINSON, H., KLOPFER, D., FELTOVITCH, P. & WANG, Y. (1988). in M. T. H. CHI, R. GLASER & M. J. FARR *The nature of expertise*. Hillsdale, NJ: Lawrence Erlbaum (pp 310-342).
- LIFE, M. A., BARBER, P. & EDWORTHY, J. (1996). Back to the future : Cognitive ergonomics six years on. *Ergonomics*, 39, 341-344.
- LIU, X. (1991). Hypotheses testing by fundamental knowledge. *International Journal of Man-Machine Studies*, 35, 409-427.
- MINSKY, M. (1986). *The Society of Mind*. New York, Simon & Schuster.
- MORAY, N. (1987). Intelligent aids, mental models, and the theory of machines. *International Journal of Man-Machine Studies*, 27, 619-629.
- NOVICK, L. R. & HOLYOACK, K. J. (1991). Mathematical problem solving by analogy. *Journal of Experimental Psychology : Learning, Memory and Cognition*, 17, 338-415.
- NOVICK, L. R. (1988). Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology : Learning, Memory and Cognition*, 14, 510-520.
- OCHANINE, D. (1978). Le rôle des images opératives dans la régulation des activités de travail. *Psychologie et Education*, 2, 63-72.
- PATRICK, J. (1993). Cognitive aspects of fault-finding training and transfer. *Le Travail Humain*, 56, 187-209.
- PERRUCHET, P. (1997). Le rôle de l'attention dans les apprentissages implicites. In *Attention et contrôle cognitif*. Rouen : Publications de l'Université de Rouen (pp. 19-27).
- RANDEL, J. M. & PUGH, H. L. (1996). Differences in expert and novice situation awareness in naturalistic decision making. *International Journal of Human-Computer Studies*, 45, 579-587.
- RASMUSSEN, J. & JENSEN, A. (1974). Mental procedures in real life tasks. A case study in electronics trouble shooting. *Ergonomics*, 17, 293-307.
- RASMUSSEN, J. (1986). *Information processing and human-machine interaction*. North Holland : Elsevier Science.
- RAUTERBERG, M. (1995). About faults, errors and other dangerous things. In H. Stassen & P. Wieringa (Eds) *Proceedings of the XIV European annual conference on human decision making and manual control*, Delft University of Technology (session 3-4, pp. 1-7).

- REASON, J. (1987a). A preliminary classification of mistakes. in J. RASMUSSEN, K. DUNCAN & J. LEPLAT. (Eds) *New technology and human error*. Chichester, UK : Wiley & Sons Ltd.
- REASON, J. (1987b). Cognitive aids in process environments: Prostheses or tools? *International Journal of Man-Machine Studies*, 27, 463-470.
- REASON, J. (1990). *Human error*. Cambridge: Cambridge University Press.
- REASON, J. (1995). A systems approach to organized error. *Ergonomics*, 38, 1708-1721.
- REASON, J. (1997). *Managing the risks of organisational accidents*. Ashgate, Aldershot.
- REASON, J. (2000). Human errors. models and management. *British Journal of Management*, 320, 768-770
- ROEDIGER, H. L. (1980). Memory metaphors in cognitive Psychology. *Memory and Cognition*, 8, 231-246.
- SALLMINEN, S. & TALLBERG, T. (1996). Human errors in fatal and serious occupational accidents in Finland. *Ergonomics*, 39, 980-988.
- SCHANTEAU, J. (1992). Competence in experts : The role of tasks characteristics. *Organisational Behavior and Human Decision Processes*, 53, 252-266.
- SCHRAAGEN, J. M. & SCHAAFSTAL, A. M. (1996). Training of systematic diagnosis : A case study in electronics troubleshooting. *Le Travail Humain*, 59, 5-21.
- SHANK, R. & ABELSON, R. (1977). *Scripts, plans, goals and understanding*. Hillsdale, NJ, Erlbaum.
- SIMON, H. A. (1957). *Models of Man*. New York, Wiley.
- SMITH M. U. (1992). Expertise and the organization of knowledge : unexpected differences among genetic counselors, faculty, and students on problem categorization tasks. *Journal of Research in Science Teaching*, 29, 179-205.
- SPERANDIO, J.-C. (1995). L'ergonomie cognitive : améliorer les aspects cognitifs du travail humain. *Psychologie Française*, 40, 3-11.
- SUCHMAN, L. A. (1987). *Plans and situated actions. The problem of human-machine communication*. Cambridge University Press, Cambridge.
- VAN ELSLANDE, P. (1992). Les erreurs d'interprétation en conduite automobile : mauvaise catégorisation ou activation erronée de schémas ? *Intellectica*, 15, 125-149.
- ZAJCHOWSKI, R. & MARTIN, J. (1993). Differences in the problem-solving of stronger and weaker novices in physics : knowledge, strategies, or knowledge structure? *Journal of Research in Science Teaching*, 30, 459-470.

10 ACKNOWLEDGEMENTS

This paper was written at the University of Newcastle upon Tyne within the DIRC project (<http://www.dirc.org.uk>), a UK-based interdisciplinary research collaboration on the dependability of computer-based systems. The authors wish to thank David Greathead (University of Newcastle upon Tyne, UK), Peter Popov (City University, UK), Michael Hildebrandt (University of York, UK), Mark Rouncefield (University of Lancaster, UK) and anonymous reviewers for useful comments. The authors are also grateful to EPSRC for funding this research.