Pergamon

Control Eng. Practice, Vol. 5, No. 3, pp. 407-416, 1997
Copyright © 1997 Elsevier Science Ltd
Printed in Great Britain. All rights reserved
0967-0661/97 $17.00 + 0.00

PII:S0967-0661(97)00018-X

OPERATOR ADAPTATION IN PROCESS CONTROL: A THREE-YEAR RESEARCH PROGRAM

K.J. Vicente

Department of Mechanical and Industrial Engineering, Cognitive Engineering Laboratory, University of Toronto,
5 King's College Road, Toronto, Ontario, Canada M5S 3G8 (benfica@mie.utoronto.ca)

(Received September 1996; in final form January 1997)

Abstract: This paper provides an integrative summary of a three-year research program investigating various factors pertinent to human operator adaptation in process-control systems. Four longitudinal experiments were conducted with a simplified but representative thermal-hydraulic process simulation. These experiments investigated the impact of four behaviour-shaping constraints that can influence operator adaptation: interface content, interface form, type of training, and pre-existing competencies. The findings obtained from the research program are summarized, and a number of implications for the design and operation of process-control systems are suggested.

Keywords: Human-machine interfaces, cognitive systems, fault detection, fault diagnosis, graphic displays, human factors, nuclear power stations, process control

1. INTRODUCTION

The greatest threat to the safety of process control systems, such as nuclear power plants (NPPs), is events that are not familiar to operators and that have not been anticipated by designers (Vicente and Rasmussen, 1992). Accident reports consistently reveal that it is precisely under these conditions that large-scale disasters occur. Under these challenging circumstances, the operator's role is one of adaptive problem solver. Because the event has not been anticipated by system designers, the available procedures, experience, and automated aids are not directly applicable. Consequently, operators must improvise a solution based on their understanding of the functional structure of the plant. This is a very demanding task, and the potential consequences of failure are enormous. Therefore, instead of *expecting* operators to effectively adapt to such unanticipated contingencies on their own, it is important to *systematically design* human-machine interfaces that explicitly support operators in effectively adapting to system demands.

There are already some efforts underway in industry to achieve this goal. For example, Toshiba in Japan have undertaken a large-scale research and development effort directed toward designing advanced control rooms for a next-generation NPP (Itoh, *et al.*, 1995). The ultimate goal of this project is to ensure that future control rooms effectively support operator adaptation under all operating conditions, particularly unanticipated events. This effort is notable because a prototype advanced control room has been implemented on a very large scale, namely a full-scope boiling water reactor simulator. Although a great deal of thought and effort has gone into the design of this prototype, it has yet to be empirically evaluated. Thus, it is not yet known how well the Toshiba design supports operator adaptation.

This limitation is not specific to this particular project, but rather reflects a gap in the research literature. Very little systematic, controlled, empirical research has been conducted on human operator adaptation in process control systems (for exceptions, see (Crossman and Cooke, 1962/1974) and (Moray, *et al.*, 1986)). It is important that this gap be reduced so that control room designers in industry have some guidance in making decisions in the design of future plants. Otherwise, unanticipated events will continue to lead to large-scale disasters that have enormous economic costs and that threaten human life and the natural environment.

This paper addresses this issue by presenting the results of a three-year research program investigating operator adaptation in process-control systems. First, a conceptual framework for investigating human adaptation in process-control systems will be presented. Second, the four studies comprising the research program will be briefly described. Third, the findings of these studies will be integrated using the framework presented earlier. Fourth and finally, the design implications and other contributions of the research program will be summarized. These should be of particular interest to engineers in industry who are responsible for designing advanced control rooms for process-control systems.

2. CONCEPTUAL FRAMEWORK

Figure 1 shows some of the forces that can impact human adaptation in complex work environments. Four factors are listed: the information content represented in the interface; the visual form of the interface; the type of training operators receive; and, the competencies that operators already bring to the work situation. (Social-organizational factors are also crucial but they are outside of the scope of this work.)

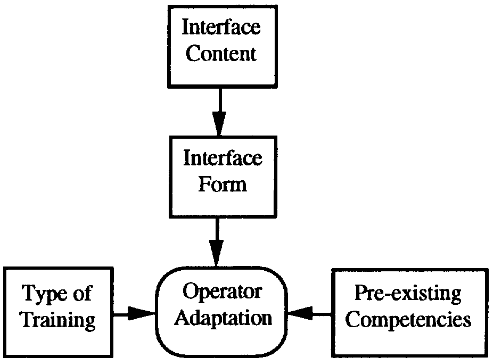

Fig. 1. Conceptual framework for the research program. Four factors shape operator adaptation: interface content, interface form, type of training, and pre-existing competencies.

2.1 Interface content

One of the strong constraints on adaptation is the information presented in the interface. If irrelevant information is presented, then adaptation will almost certainly be hampered since operators will have to determine what information to attend to, and what information to ignore. In the absence of relevant conceptual knowledge, this discrimination task can only proceed inductively on a trial-and-error basis, and is therefore likely to be slow.

Conversely, if critical information is not presented in the interface, then adaptation should also be impaired but for different reasons. If there are constraints between system variables which are not explicitly represented, then it will be very difficult for operators to take those constraints into account in controlling the system. As the findings of Sanderson, *et al.* (1989) show, it is difficult to adapt to what is not displayed. Operators will thereby have to infer the missing constraints based on observation of other system variables.

A good example of the effect of interface content on adaptation is provided by Moray, *et al.* (1986). In the process simulation that subjects were controlling in that experiment, there was a uni-directional coupling between the two system goal variables, temperature and volume. Changing volume by manipulating the mass flow could affect temperature, but changing temperature by manipulating the heater would not affect volume. Given this constraint, it would be better for subjects to stabilize volume at the desired setpoint first, and only then worry about satisfying the temperature goal. Doing the inverse would be inefficient since changing volume second would change the temperature that had just been achieved. However, because the uni-directional coupling between temperature and volume was not immediately visible in the interface, some subjects initially adopted the inefficient strategy. In fact, one subject persisted in doing so after about 12 hours of practice. Explicitly representing the uni-directional constraint in the interface would likely help subjects by making them aware that such a constraint exists. This example illustrates the importance of considering interface content as a contributing factor in operator adaptation.

2.2 Interface form

Even if two interfaces are equivalent in terms of content, performance can also vary as a function of visual form. As operators gain more and more experience with a system, they become increasingly attuned to the perceptual properties of the interface (Rasmussen, 1986). In this way, the interface can be viewed as manipulating the operators' attention. This has important implications for human adaptation.

For example, if the interface is designed in such a way that it attracts operators' attention to interface attributes which are either irrelevant or not critical, then this may make it difficult for operators to attend to the most informative aspects of the interface. The result would likely be dysfunctional adaptation, thereby leading to poor performance. In contrast, if the interface is designed in such a way that the perceptual saliency of its objects is commensurate with their relative importance, then this might facilitate functional adaptation. The interface would serve as a guide, distributing the operators' attention over display objects in a meaningful way.

MacGregor and Slovic (1986) showed the impact that interface form can have. They presented subjects with two different displays with the same information content. However, in one display, the most important variables were mapped onto the most perceptually salient features of the display and the least important variables to the less salient features. In the other, there was an arbitrary mapping between the importance of each variable and the salience of the feature that was used to represent that variable. Results indicated that the commensurate mapping led to better performance, thereby indicating the importance of using interface form to shape adaptation in a functional direction.

2.3 Type of training

The type of training that operators receive also influences their adaptation processes. First, initial training provides operators with a baseline set of job competencies. Second, training also exerts a top-down influence on adaptation in that it provides operators with some guidance on what to look for and what is considered important in controlling the system of interest. Third, training can also assist operators to deal with unanticipated faults. In these cases, the operators' experience base will not be of direct help. However, principles that were taught during training are a resource that may allow operators to go beyond their experience base and thereby effectively deal with novel events (cf. (Patrick and Haines, 1988)).

A topic that has been relatively ignored is the interactive effect between training and interface design on operator adaptation. Many training studies in process control have produced findings that have led some to argue that theoretical instruction about the principles governing process behavior is of no help or even a hindrance to operators (e.g., (Crossman and Cooke, 1962/1974)). However, as Lees (1974) points out, this conclusion is surely incorrect. A more cautious inference is that the conditions under which theoretical training can have a positive impact on performance have yet to be determined. One untested hypothesis is that the reason theoretical training has often failed to improve performance is because the interface was not designed to reflect the theoretical principles governing the process. To take a concrete example from the control of water-based power plants, what good is it instructing operators about the Rankine Cycle if the information they receive from the interface is not mapped onto that cycle? If one were to base *both* training and interface design on fundamental physical principles, then perhaps one could show the benefits of theoretical instruction. This suggests that training can be an important factor in the investigation of human adaptation.

2.4 Pre-existing competencies

The competencies that operators bring to the job can also have a marked effect on the adaptation processes they exhibit. Such pre-existing competencies can take many forms, such as cognitive style, declarative knowledge, procedural knowledge, perceptual and motor skills, and population stereotypes. Although this factor is difficult to control for (except perhaps through selection), it is important that it be explicitly taken into account in experimentation since it can mediate the effects of other more controllable influences on adaptation. Thus, one of the most important reasons for considering pre-existing competencies is that they can help explain why different people often exhibit very different patterns of behavior under what seem to be identical experimental conditions. In fact, the situations are not the same in the sense that each person comes to the experiment with a different set of competencies. These competencies, which can vary considerably across individuals, can be thought of as the initial conditions for the adaptation process.

Dörner (1989) provides several examples that point to the importance of pre-existing competencies. In his Lohhausen experiment, subjects were asked to act as the mayor of a computer-simulated town and to look after the well-being of its inhabitants. The influence of subjects' backgrounds on their behavior was marked: a subject who is a school teacher focused his activities on particular teaching problems in the town school system and left the rest of society to its own; a left-wing subject decided to introduce extensive socialistic measures and attributed difficulties to reactionary sabotage; and finally, a student subject without any experience in production planning used a very simple analogy from her daily life as a basis for her decision making. These examples indicate that knowledge of pre-existing competencies is essential to understanding the idiosyncratic adaptation processes exhibited by individuals.

2.5 Summary

Four factors have been identified as being relevant to the study of human adaptation: interface content, interface form, type of training, and pre-existing competencies. The first three of these can be systematically manipulated, whereas the last category

Table 1. Integrated summary of the three-year research program, showing the behaviour-shaping constraints investigated in each experiment. See text for a more detailed description.

	Behaviour-Shaping Constraints			
Experiment #	Interface Content	Interface Form	Type of Training	Pre-Existing Competencies
1	P vs. P+F	P vs. P+F	None	Demographic Data + Cognitive Style
2	P vs. P+F	P vs. P+F	None vs. AH	
3	P+F	P+F vs. P+F/Div.	None	
4	P+F	P+F	None vs. Review vs. Review + SE	

needs to be taken into account if differences in adaptation across subjects are to be well understood. All of these factors have been incorporated into the research program described next.

3. RESEARCH PROGRAM

In this section, the interface design principles, research vehicle, and experiments that comprised this research program are briefly summarized. This will provide the necessary background for the following section which will summarize the findings obtained during this programmatic series of studies.

3.1 Interface design principles

The manipulations of interface design content and form conducted in this research program have been motivated by the principles of ecological interface design (EID), a theoretical framework for designing interfaces for complex work environments (Vicente and Rasmussen, 1990; 1992). EID is based on the skills, rules, and knowledge taxonomy of levels of cognitive control (Rasmussen, 1986). The framework includes three prescriptive design principles, each directed at providing the appropriate interface support for a specific level of cognitive control. First, to support skill-based behaviour, operators should be able to act directly on the interface. Second, to support rule-based behaviour, the interface should maintain a consistent one-to-one mapping between the work domain constraints and the perceptual cues provided in the interface. Third, to support knowledge-based behaviour, the interface should represent the work domain in the form of an abstraction hierarchy (AH) (Rasmussen, 1986), which can serve as an externalized mental model to support problem solving. This system model contains both physical and functional representations of the system. Vicente and Rasmussen (1990, 1992) provide a more detailed description and justification of these design principles.

3.2 Research vehicle

The experiments comprising this research program were all conducted within the context of a single, carefully-chosen research vehicle, DURESS (DUal REservoir System Simulation) II (Pawlak and Vicente, 1996). DURESS II is a real-time interactive thermal-hydraulic process-control simulation consisting of several pumps, valves, heaters, and reservoirs. The simulation can be configured to present startup, tuning, and shutdown tasks, as well as various types of faults, including multiple, simultaneous faults.

3.3 Representative design

DURESS II was explicitly designed to be *representative* (Brunswik, 1956) of complex work domains, albeit on a much smaller scale (Vicente, 1991). For example, the components in DURESS II are governed by first-order lag dynamics, thereby introducing a time lag between subjects' actions and the resulting impact on the process. Furthermore, DURESS II is composed of several interacting subsystems, thereby introducing an element of structural complexity. Also, the simulation includes an element of risk by modeling equipment failures that are caused by subjects' dysfunctional actions (e.g., heating an empty reservoir). Conducting this research program with a representative research vehicle like DURESS II increases the chances that the research results obtained will generalize to the industrial-scale process-control systems of interest.

Representative design was also adopted in a number of other ways in this research program:

1. Subjects were given extensive experience at controlling the system so that the results would be more likely to generalize to experienced operators.
2. A range of representative faults were introduced into the simulation, including both single- and multiple-fault scenarios.

3. Faults occurred relatively infrequently, just as in industrial systems.
4. Subjects were not told what faults might occur during the experiment, when they would occur, or how frequently they would occur. Thus, faults occurred unexpectedly, just as in industrial systems.

Although laboratory research is rarely conducted under such representative conditions, these characteristics are essential if results are to generalize outside of the laboratory to industrial process-control systems.

3.4 Experiments

These general methodological details provide the background for discussing the experiments that comprised this research program. Table 1 shows the relationship between the factors in the conceptual framework described in the previous section and the four experiments. In all studies, subjects controlled DURESS II during each week day for approximately 1 hour. In experiment #1, this quasi-daily practice lasted for an unprecedented 6 months, whereas in the remaining three studies it lasted for approximately 1 month. The longitudinal nature of these experiments thereby provided an opportunity to observe subjects' adaptation processes. Note that data for two types of pre-existing competencies (i.e., demographics and cognitive style) were collected in all four studies (in some cases, after the fact).

Experiment #1. As indicated in Table 1, the focus of experiment #1 was to assess the impact of interface content and form on long-term adaptation (Christoffersen, *et al.*, 1994). Two interfaces for DURESS II were compared, a traditional interface that contained a physical (P) representation of the system and an interface based on the principles of EID which contained both physical and functional (P+F) representations of the system (see (Vicente and Rasmussen, 1990) for descriptions of these interfaces and the process by which the P+F interface was designed according to the principles of EID). It was expected that the P+F interface would lead to better performance, primarily under fault conditions. Note that these two interfaces differed in terms of both content (the P+F contained all levels of the AH, whereas the P interface only contained a subset of those levels) and form (the P+F interface contains emergent features whereas the P interface consists primarily of component icons).

Experiment #2. As shown in Table 1, the objective of experiment #2 was to investigate the interaction between interface design and model-based training on adaptation (Hunter, *et al.*, 1995). A 2 x 2 design was adopted with two levels of interface design (P vs. P+F) and two levels of training (none vs. AH). No study had ever investigated the impact of training based on an AH representation of a system. The hypothesis was that there would be an interaction between these two factors such that the combination of training and the P+F interface would lead to a synergistic improvement in performance. The rationale for this prediction is that there is benefit to be gained from basing both training and interface design on a common model of the system (i.e., the AH). In addition, it was expected that the P+F interface and the training based on the AH would each individually lead to better performance.

Experiment #3. As illustrated in Table 1, the focus of experiment #3 was to investigate the impact of interface form (Howie, *et al.*, 1996). More specifically, the P+F interface for DURESS II, based on the principles of EID, was compared with a divided P+F interface consisting of four different windows, each representing one of the levels of the AH displayed in the P+F interface. Both interfaces contain the same information at four levels of the AH. The primary difference between them is that the integrated P+F interface presents these levels of abstraction in a single, integrated display, whereas the divided P+F interface presents these levels on four separate displays which must be accessed serially (i.e., only one level of abstraction could be viewed at any one time, but subjects could move freely among levels). It was expected that the divided P+F interface would impair performance because it causes a deterioration in visual momentum (Woods, 1984). The strategies that subjects used to navigate between levels of abstraction in the divided P+F interface were also of relevance. An additional goal was to understand which navigation strategies led to good and poor performance. Furthermore, the good strategies might shed some light on how an AH representation should be parsed onto multiple display pages for large-scale systems where the entire system is too large to be represented in one display (unlike DURESS II).

Experiment #4. As shown in Table 1, the objective of experiment #4 was to investigate a second type of training (Howie, *et al.*, 1996). In particular, the effect of instructing subjects to review and self-explain their control performance was evaluated. Previous work in cognitive science had shown that proficient students spontaneously engaged in self-explanation more frequently than poor students (Chi, *et al.*, 1989). A subsequent study showed that instructing students to engage in self-explanation improved the performance of both proficient and poor students equally (Chi, *et al.*, 1994). It was thought that the same result might apply to the domain of process control. That is, subjects who are proficient controllers of DURESS II would be more likely to spontaneously engage in self-explanation than poor subjects. Moreover, instructing subjects to engage in self-explanation should improve their performance with the P+F interface. These hypotheses were put to the test by comparing the performance of three groups of subjects. The first group controlled DURESS II with the P+F interface but did not review their performance and were not instructed to self-explain their control performance. A second group also used the P+F interface but periodically reviewed their performance on the trial they had just completed. This was accomplished by using a computer program that replayed subjects' actions and the evolution of the state of the system on the same interface subjects used during the trial. Finally, the third group also used the P+F interface and periodically reviewed their performance in the manner just described, but was also explicitly encouraged to self-explain their control performance while reviewing their trials.

4. SUMMARY OF RESEARCH FINDINGS

This section provides a brief summary of the findings obtained over the integrated, three-year research program. A more detailed description of the results can be found in the reports associated with each phase of the program (Christoffersen, *et al.*, 1994; Hunter, *et al.*, 1995; Howie, *et al.*, 1996). The findings are organized according to the four factors identified by the conceptual framework in Figure 1.

4.1 Interface content

The impact of interface content on adaptation was investigated in experiments #1 and #2 by comparing performance with the P and P+F interfaces for DURESS II. The latter is based on an AH analysis and therefore contains both physical and functional information, whereas the P interface only displays the subset of the AH pertaining to physical information. Note that there are also differences in form between these two interfaces, so it is not possible to determine how much of the effects described in this section are due to form vs. content, although a previous study established that the performance advantage of the P+F interface is not due to differences in form alone (Vicente, 1992a). Clearly, both factors are important since no form can make up for missing content and even the most relevant content is useless if it is displayed in a form that is not legible. With this limitation in mind, the impact of this content manipulation on normal trials, fault management trials, and deep knowledge will be discussed next.

The results from experiment #1 revealed very little difference in the average performance of the two interface groups on normal trials. There were some indications that the P+F interface led to slower performance initially, which is not surprising, given that it is more visually complex than the P interface. The other significant effect indicated that the P interface led to faster performance than the P+F interface during shutdown tasks. Results from the transfer trials indicate that this may be a subject effect rather than an interface effect. All other comparisons between interfaces failed to show any significant differences in average performance. This may have resulted from the fact that there were only 3 subjects in each group. This interpretation is supported by the results of experiment #2 which showed statistically significant faster performance for the P+F group on normal trials, even after extensive practice. Note that experiment #2 was conducted under more stringent conditions, since the tolerance on the temperature goal was only ±1.5 °C compared to ±2 °C in experiment #1. Taken together, these results suggest two conclusions for normal trials. When tighter control is required (as in experiment #2), presenting all levels of the AH leads to faster average performance than presenting only physical information. When the control requirements are not as tight, there does not seem to be any substantial performance cost to adding the functional levels of information in the AH to an interface, at least once operators are experienced.

The effect of interface content on performance for normal trials is more noticeable for performance variability. The results of experiment #1 indicate that the P group exhibited significantly less consistent performance than the P+F interface, occasionally taking up to 2 times longer than usual to complete the required tasks, even after 5.5 months of quasi-daily practice. A transfer manipulation indicated that this effect was specific to the interface, not to the subjects. The improved consistency induced by the P+F interface was corroborated in experiment #2. These results indicate that adding functional levels of information defined by the AH significantly reduces the variability in performance.

As for fault-management performance, experiments #1 and #2 painted a consistent picture. Subjects with the P+F interface usually exhibited faster detection times, more accurate diagnoses, and faster compensation times than subjects with the P interface. Thus, the advantages of representing all levels of the AH in an interface are more prominent under fault conditions than they are under normal conditions, as expected.

Interestingly, analyses of verbal protocols indicate that these performance differences are linked with differences in strategies between groups. Subjects with the P interface frequently only detected faults after acting on the system several times in succession, and then eventually observing that their actions were not having the usual effect, indicating the presence of a problem. In other cases, P subjects detected faults by noticing that the rules that they usually used to control the system were not having their desired effect. There are two disadvantages to these strategies. First, they generally lead to slow fault detection times. Second, they give very little information beyond the fact that there is a fault somewhere in the system. In contrast, the P+F subjects were more likely to detect faults by directly observing the interface and noticing that the constraints that usually governed the system had been violated. This strategy is usually faster than the other two. Even more importantly, it allows subjects to localize the fault to a specific part of the system, thereby facilitating fault diagnosis and compensation. These strategies are consistent with the theoretical motivation behind using the AH as a basis for determining interface content (Vicente and Rasmussen, 1992).

These results are reinforced by those of experiment #3, which clearly indicated that effective performance is associated with the access of functional information. In that study, subjects who consulted functional information less frequently tended to be outperformed by subjects who consulted functional information more frequently. This finding shows the importance of representing the functional levels of the AH, something which traditional interfaces do not do comprehensively.

Finally, the longitudinal results of experiment #1 also lead to inferences about the impact of including the AH in an interface on subjects' deep knowledge. The results obtained from knowledge elicitation tests revealed that the most proficient P subject thought about the system in terms of a sequence of quantitative actions, much like one would find in a detailed set of procedures. Although this "script" generally worked well when there were no faults, as the subject himself pointed out, he did not really understand why they were effective. In contrast, the most proficient P+F subject thought about the system in terms of the goals or functions that needed to be satisfied and did not focus very much on memorizing a sequence of actions. This is precisely the kind of thinking that the P+F interface is intended to engender (Vicente and Rasmussen, 1992). Because the interface provides a visualization of the system constraints, one does not need to rely as much on procedures. Rather, one can make more use of the feedback being provided by the interface as an error signal for controlling the process. Note that the sample size in this experiment was unusually low, so these results should be interpreted with caution. With this in mind, the results from this 6-month study indicate that including physical and functional levels of the AH in an interface can induce some subjects to acquire a deeper understanding of the system than an interface with physical information alone.

4.2 Interface form

Two experiments bear on the effect of visual form on adaptation. Experiment #1 investigated adaptation over a 6 month period, thereby evaluating the possibility that the emergent feature graphical displays in the P+F interface might be used by subjects as a cognitive crutch, leading to a deskilling effect. Experiment #3 investigated the impact of displaying each level of the AH in a separate visual form that could only be accessed serially.

Two phenomena observed in experiment #1 are relevant to the issue of interface form. The first was a perceptual fixation effect occurring during a reservoir leak fault that had been observed in a previous study with the P+F interface (Pawlak and Vicente, 1996). Subjects were accustomed to making the emergent feature in the mass balance display (i.e., a sloped line indicating the intended rate of change of volume) vertical to stabilize volume. However, in the presence of a leak, this sloped line must have a positive slope to offset the leak. Some subjects insisted on trying to make the line vertical several times in succession, even though each time they did so, the volume would decrease. This suggests that emergent feature displays, like the P+F interface for DURESS II, may lead to perceptual fixation effects. However, subsequent trials indicated that the fixation effect disappears after subjects have observed the behaviour of the interface for one reservoir leak fault. Thus, it is perhaps more accurate to say that this fixation effect is not so much a defect with the display as it is a part of learning how to use the display under a wide variety of conditions. (See (Pawlak and Vicente, 1996) for a more detailed discussion of this point.)

The second phenomenon pertinent to interface form observed in experiment #1 is that the least proficient subject in the P+F group exhibited a very shallow understanding of the system. He relied heavily on the perceptual features of the display to control the system, rather than thinking deeply about the meaning of those features. As a result, his performance was particularly bad on fault trials. Moreover, when he transferred to the P interface, after a few trials he was no longer able to stabilize the system because the graphics that he relied upon were not available anymore. This result suggests that emergent feature graphical displays can lead to a form of deskilling when subjects do not reflect upon the meaning of the information provided to them. Nevertheless, it is important to point out that the performance of this subject was no worse than that of the least proficient P subject, who also did not think deeply about the system. Thus, it seems that when subjects do not reflect on their activities, then performance is poor, even with a traditional interface that does not have emergent feature graphical displays. Thus, traditional interface design practices do not provide a viable alternative to EID.

Experiment #3 investigated the impact of displaying the information in the P+F interface in four visual forms that had to be viewed serially, rather than one integrated form that made all information available in parallel. Not surprisingly, forcing subjects to search for information across windows significantly impaired performance. This study also led to implications for how to parse an interface based on an AH representation into multiple visual forms (when the system is too large to be displayed in one integrated form). The "flows" display was consulted infrequently by all subjects. The reason for this seems to be that flow information is redundant with that provided in the settings display under normal conditions. Thus, a knowledgeable observer who understands the relationships between settings and flows need only consult the settings display. However, during faults, this is no longer true since the constraints that govern the relationship between flows and settings may be violated by the fault. As a result, it is important to consult both levels to detect that a fault has occurred (i.e., that a constraint has been violated) and to diagnose the fault. The results from experiment #3 are consistent with this observation. Subjects who never, or very rarely, consulted the flows display performed poorly compared to subjects who made a point of periodically checking the flow information. These findings suggest that it is important to try to parse the AH representation into visual forms that do not cut across process constraints. Doing so should allow operators to detect constraint violations (i.e., faults) within a display, rather than having to integrate information across displays.

4.3 Type of training

The impact of two different forms of instruction were investigated in this research program. Experiment #2 evaluated the impact of training based on an AH representation of the process. Experiment #4 evaluated the impact of instructing subjects to review their control performance retrospectively and to self-explain their actions while doing so.

The results of experiment #2 showed that training based on an AH representation of DURESS II resulted in a larger improvement in performance than merely giving subjects practice at controlling the system. This greater increase in performance was observed on normal and fault trials. Moreover, a content analysis of subjects' retrospective comments after fault trials indicates that subjects who received the training exhibited a deeper understanding of the system. These results indicate that training based on the AH can improve performance. It might seem self-evident that theoretical training is better than no training, but very few studies in the literature had shown this effect, and none had done so using training based on an AH analysis. Contrary to expectations, there was virtually no indication of an interaction between training and interface. Thus, the P and P+F groups seemed to benefit equally from the AH training.

Experiment #4 was motivated by experiment #1 which suggested that subjects who reflected on the feedback in the P+F interface achieved a more proficient level of performance than those who seemed to adopt a shallow approach to learning. This implied that instructing subjects to retrospectively review and self-explain their control performance might help them make the most of the P+F interface. The results indicated no difference between groups on normal trials. For fault trials, self-explanation improved subjects' performance but this advantage was not statistically significant. Interestingly, when the self-explanation subjects were divided at the median into proficient and poor performers, analyses revealed that more proficient subjects engaged in self-explanation more frequently. Furthermore, proficient subjects explained the rationale for their actions more frequently than the less proficient subjects. These results suggest that merely instructing subjects to self-explain may not improve performance, but that subjects who have a greater propensity for spontaneously engaging in self-explanation tend to be more proficient performers with an EID interface.

4.4 Pre-existing competencies

In each study, demographic data were obtained for all subjects. The only factor which had a consistent relationship with performance was whether or not subjects were in the prestigious Engineering Science undergraduate program at the University of Toronto. These students go through a very rigorous selection process, and so it is not surprising to find that they were proficient performers. This effect may be due to a higher general level of intelligence, or to a specific set of knowledge or attributes that lead to more effective performance.

Part-way through the research program, it was observed that the holist-serialist cognitive style (Pask and Scott, 1972) might account for some of the individual differences between subjects. From that point onward, an attempt was made to bring back subjects who had participated in previous studies to evaluate their cognitive style. This effort was relatively successful since data were obtained for most of the subjects who had participated in experiments #1 and #2. A preliminary analysis of the data indicates that there is no significant relationship between cognitive style and performance with the P interface. With the P+F interface, however, there is a significant positive correlation between performance and degree of holism (Howie, 1996). If these results are confirmed with a larger sample and with more sophisticated analyses, then this would suggest that operators should be selected so that they have a holist cognitive style if the full benefits of EID are to be realized.

5. CONCLUSIONS

The three-year research program described in this paper adds to the existing meagre understanding of operator adaptation in process-control systems. With a small number of experiments, it is very difficult to derive definitive conclusions. Furthermore, since the studies were conducted with a small-scale simulation, it is not known to what extent the results generalize to industrial systems (although the attention paid to representative design increases the likelihood of generalization). Because of these limitations, it would be premature to use the results of this research program to make strong recommendations for the design of advanced control rooms for future plants.

Nevertheless, from a more pragmatic perspective, it is important to point out that designers in industry are currently designing advanced computer-based control rooms for next-generation plants (Vicente, 1992b). In the absence of pertinent research, designers can only rely on their subjective opinions. Thus, the research presented here can serve a useful practical role by helping designers focus their analysis, design, and evaluation activities on issues that seem to be particularly promising or critical. The following design implications are put forth in this spirit:

1. EID is a promising framework for designing interfaces for process-control systems, since it has demonstrable advantages over more-traditional interface design approaches. However, having an interface based on the principles of EID does not guarantee effective performance.
2. To experience the benefits of EID, it seems likely that operators need to be trained to think functionally rather than procedurally. This may require a fundamental shift in process-control plant operation philosophy.
3. Similarly, there is some preliminary evidence to indicate that, to experience the benefits of EID, operators should be selected to have a holistic rather than a serialist cognitive style.
4. Training based on an AH representation of a system can improve operator performance.
5. Not surprisingly, an EID interface that integrates levels of the AH into a single display will lead to better performance than an EID interface that only provides serial access to each level of abstraction individually.
6. Operators' information search patterns will be critical in determining the success of an interface based on EID. More specifically, operators who consult functional levels of information more frequently will likely outperform operators who consult such levels less frequently.
7. It is important to try to parse the AH representation into visual forms that do not cut across process constraints. Doing so should allow operators to detect constraint violations (i.e., faults) within a display, rather than have to integrate information across displays.
8. Requiring operators to review their performance and to self-explain their actions will not greatly improve the level of performance observed with an interface based on EID. However, operators who spontaneously self-explain more frequently and more thoroughly will probably outperform operators who self-explain less frequently and less thoroughly with an interface based on EID.

These suggestions need to be followed up by full-scope simulator studies investigating operator adaptation, in order to make users confident that future control rooms will lead to improved productivity and safety.

ACKNOWLEDGEMENTS

This research was sponsored by contracts from the Japan Atomic Energy Research Institute (Dr. Fumiya Tanabe, Contract Monitor), and by grants from the Natural Sciences and Engineering Research Council of Canada. The author is indebted to all of the people who contributed to this research, especially: Klaus Christoffersen, Dianne Howie, Chris Hunter, Michael Janzen, Darryl Minard, Dimitrios Nastos, Dr. John Patrick, Dr. Tanabe, and the 54 participants who devoted so much time and effort to these experiments.

REFERENCES

Brunswik, E. (1956). *Perception and the representative design of psychological experiments*, (2nd ed.). Berkeley, CA: University of California Press.

Chi, M. T. H., Bassok, M., Lewis, M. W., Reimann, P., and Glaser, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, *13*, 145-182.

Chi, M. T. H., De Leeuw, N., Chiu, M.-H., and LaVancher, C. (1994). Eliciting self-explanations improves understanding. *Cognitive Science*, *18*, 439-477.

Christoffersen, K., Hunter, C. N., and Vicente, K. J. (1994). *Research on factors influencing human cognitive behaviour (I)* (CEL 94-05). Toronto: University of Toronto, Cognitive Engineering Laboratory.

Crossman, E. R. F. W., and Cooke, J. E. (1962/1974). Manual control of slow-response systems. In E. Edwards and F. P. Lees (Eds.), *The human operator in process control* (pp. 51-66). London: Taylor and Francis.

Dörner, D. (1989). *Die logik des misslingens*. Reinbek bei Hamburg: Rowohlt.

Howie, D. E. (1996). *Shaping expertise through ecological interface design: Strategies, metacognition, and individual differences* (CEL 96-01). Toronto: University of Toronto, Cognitive Engineering Laboratory.

Howie, D. E., Janzen, M. E., and Vicente, K. J. (1996). *Research on factors influencing human cognitive behaviour (III)* (CEL 96-06). Toronto: University of Toronto, Cognitive Engineering Laboratory.

Hunter, C. N., Janzen, M., E., and Vicente, K. J. (1995). *Research on factors influencing human cognitive behaviour (II)* (CEL 95-08), Toronto: University of Toronto, Cognitive Engineering Laboratory.

Itoh, J., Sakuma, A., and Monta, K. (1995). An ecological interface for supervisory control of BWR nuclear power plants. *Control Engineering Practice*, *3*, 231-239.

Lees, F. P. (1974). Research on the process operator. In E. Edwards and F. P. Lees (Eds.), *The human operator in process control* (pp. 386-425). London: Taylor and Francis.

MacGregor, D., and Slovic, P. (1986). Graphic representation of judgmental information. *Human-Computer Interaction*, *2*, 179-200.

Moray, N., Lootsteen, P., and Pajak, J. (1986). Acquisition of process control skills. *IEEE Transactions on Systems, Man, and Cybernetics*, *SMC-16*, 497-504.

Pask, G., and Scott, B. C. (1972). Learning strategies and individual competence. *International Journal of Man-Machine Studies*, *4*, 217-253.

Patrick, J., and Haines, B. (1988). Training and transfer of fault-finding skill. *Ergonomics*, *31*, 193-210.

Pawlak, W. S., and Vicente, K. J. (1996). Inducing effective operator control through ecological interface design. *International Journal of Human-Computer Studies*, *44*, 653 - 688.

Rasmussen, J. (1986). *Information processing and human-machine interaction: An approach to cognitive engineering*. Amsterdam: North-Holland.

Sanderson, P. M., Verhage, A. G., and Fuld, R. B. (1989). State-space and verbal protocol methods for studying the human operator in process control. *Ergonomics*, *32*, 1343-1372.

Vicente, K. J. (1991). *Supporting knowledge-based behavior through ecological interface design* (Unpublished dissertation). Urbana, IL: University of Illinois at Urbana-Champaign.

Vicente, K. J. (1992a). Memory recall in a process control system: A measure of expertise and display effectiveness. *Memory and Cognition*, *20*, 356-373.

Vicente, K. J. (1992b). Multilevel interfaces for power plant control rooms I: An integrative review. *Nuclear Safety*, *33*, 381-397.

Vicente, K. J., and Rasmussen, J. (1990). The ecology of human-machine systems II: Mediating "direct perception" in complex work domains. *Ecological Psychology*, *2*, 207-250.

Vicente, K. J., and Rasmussen, J. (1992). Ecological interface design: Theoretical foundations. *IEEE Transactions on Systems, Man, and Cybernetics*, *SMC-22*, 589-606.

Woods, D. D. (1984). Visual momentum: A concept to improve the cognitive coupling of person and computer. *International Journal of Man-Machine Studies*, *21*, 229-244.