A Theoretical Framework for Ecological Interface Design

Vicente, Kim J.; Rasmussen, Jens

Publication date:
1988

Document Version
Publisher's PDF, also known as Version of record

Link back to DTU Orbit

Citation (APA):
Vicente, K. J., & Rasmussen, J. (1988). *A Theoretical Framework for Ecological Interface Design*. Risø-M No. 2736

Risø-M-2736

A Theoretical Framework for Ecological Interface Design

Kim J. Vicente and Jens Rasmussen

Risø National Laboratory, DK-4000 Roskilde, Denmark
August 1988

ISBN-87-550-1459-3

ISSN 0418-6435

Grafisk Service Riso 1988

RISØ-M-2736

A THEORETICAL FRAMEWORK FOR ECOLOGICAL INTERFACE DESIGN

Kim J. Vicente and Jens Rasmussen
Risø National Laboratory
Roskilde, Denmark

ABSTRACT. A theoretical framework for designing interfaces for complex systems is described. The framework, called ecological interface design (EID), suggests a set of principles for designing interfaces in a way that supports the fundamental properties of human cognition. The basis of EID is the skills, rules, knowledge model of cognitive control. In order to support the full range of operator problem solving activities, EID suggests how to design interfaces that simultaneously support each of the three levels of cognitive control, but that do not force processing to a higher level than the demands of the task require. The EID approach extends the concept of direct manipulation interfaces by taking into account the added complications introduced by complex systems. In this paper, we describe the development of the framework, its theoretical foundations, and examples of its application to various work domains.

1. INTRODUCTION

Recently, a new form of human-computer interaction called Direct Manipulation Interfaces (DMI) has emerged. While reports about DMI have been very enthusiastic (Shneiderman, 1983), very little is known about how to design such interfaces, or about why they are so effective (Hutchins, Hollan, and Norman, 1986). In this paper, we are concerned with extending the benefits of DMI to the area of complex systems. With this aim in mind, we begin by reviewing two existing theories of DMI, and by pointing out their limitations with respect to complex systems.

2. THEORIES OF DIRECT MANIPULATION INTERFACES

In this section, we will describe two existing theories of DMI. The first of these is the Syntactic-Semantic Model (SSM) of Shneiderman (1983), and the second is the description based on the Gulfs of Evaluation and Execution provided by Hutchins et al. (1986).

2.1. The Syntactic-Semantic Model

Shneiderman (1983) begins with the distinction between syntax and semantics. He states that syntax is an arbitrary convention, and therefore may be difficult to learn and remember. In contrast, semantic knowledge represents the user's understanding of the domain. Usually, semantic knowledge is hierarchically organized with general knowledge being decomposed into more specific concepts. Because it is meaningful, semantic knowledge is viewed as being system independent and relatively stable in long-term memory. Given this distinction, Shneiderman states that, when interacting with computers, users decompose higher level semantic concepts into lower level concepts that come closer to the syntax domain (see Figure 2a). According to this account, the advantage of DMI is that they display the objects of interest to the user, so that actions are directly in the high-level problem domain. As a consequence, the degree of decomposition required before selecting a command is reduced.

It is clear that the tight coupling between semantics and syntax is one of the keys to DMI. However, upon close examination, Shneiderman's account of the benefits of DMI is somewhat vague. Of course, this is to be expected since his paper was written at a time when the concepts of DMI were quite novel. We now recognize that there are several limitations to Shneiderman's SSM. First, the term semantics is used loosely to include the user's goals, his knowledge about how to carry out those goals (independent of syntax), and his domain knowledge. It is important to distinguish between these for they are not the same.

A second limitation of the model is that it only describes the process of forming a command (the output side of the interaction process). An equally important aspect of DMI is the processes involved in the input side, i.e., in the perception and interpretation of the displayed information. This aspect is relatively ignored, the only mention being that the objects that are displayed represent high-level semantic concepts.

Finally, the most complicated issue is exactly how the semantics should map onto the syntax. This problem is not dealt with in the model. If we accept the explanation that DMI reduce the decomposition process by making available actions that correspond to high-level domain concepts, several important questions arise. What is the appropriate 'high-level' at which the semantic concepts map onto commands? If the user has goals at multiple levels, should there be a command for each goal or only at some 'baseline' semantic level? The framework provided by the model is not robust enough to allow us to address these questions.

2.2. The Gulfs of Evaluation and Execution

Hutchins et al. (1986) also provide a theory of DMI. They begin their discussion with the premise that the general problem in human-computer interaction is that "the person's goals are expressed in terms relevant to the person - in psychological terms and the system's mechanisms and states are expressed in terms relative to it - in physical terms" (Norman, 1986, p. 38). As shown in Figure 1, the mismatch can be characterized by two gulfs between person and machine. The figure is actually a simplified version of Rasmussen's (1974; 1986) decision ladder. The Gulf of Execution refers to the gap between the person's goals and intentions, and the inputs that the computer recognizes. The Gulf of Evaluation, on the other hand, refers to the gap between the computer's output and the person's conceptual model of the task. Either of these gulfs can be bridged by the computer or by the person. Of course, placing the majority of the burden of bridging the gulfs on the person greatly increases the cognitive demands of the task, or the distance introduced by the interface (Hutchins et al., 1986). Each of the gulfs has two types of distance associated with it. Semantic distance refers to the disparity between the user's intentions and the meaning in the interface language, while articulatory distance refers to the distance between the physical form of the interface language and its meaning (Hutchins et al., 1986).

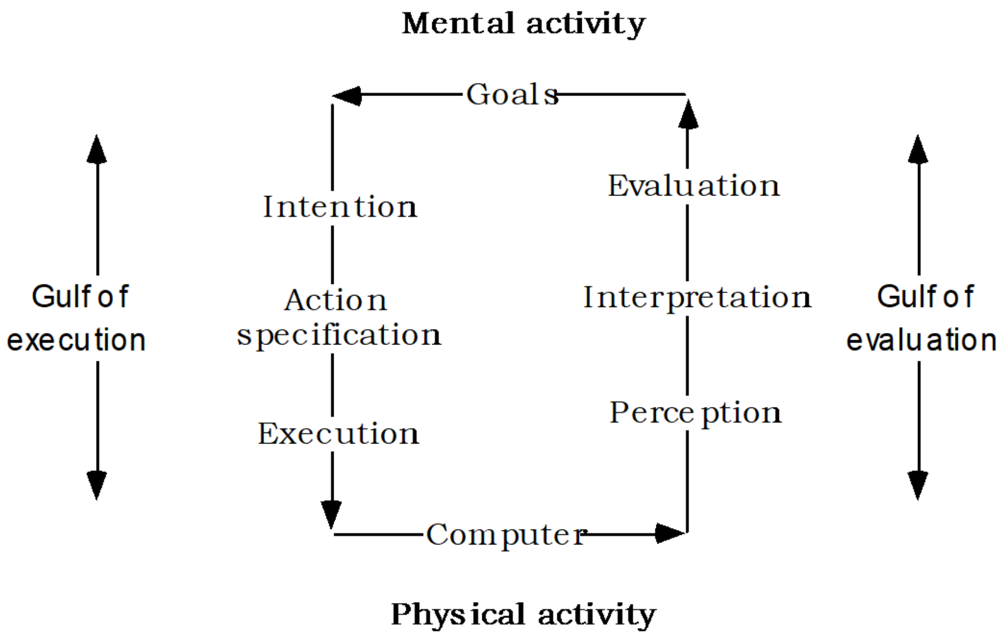

Figure 1. The gulfs of execution an evaluation. Adapted from Norman (1986).

We will first discuss the concept of semantic distance. On the execution side, semantic distance of the interface is reduced if the user's intentions can be expressed in the command language in a concise manner. The goal should be to match the level of description required by the interface language to the level at which the person thinks of the task. On the evaluation side, semantic distance is reduced if the displayed objects represent the higher level concepts that people naturally adopt when reasoning in the problem domain. In this case, the goal should be to provide a powerful, productive way of thinking about the domain.

The articulatory distance introduced by an interface is also an important factor. On the execution side, the articulatory distance is related to how closely the form of the action is to the meaning of the action. For instance, articulatory distance is reduced if the user can drag a mouse to move an object on the screen, rather than enter a command on the keyboard. On the evaluation side, articulatory distance refers to how closely the form of the displayed objects resembles their meaning. As an example, effective use of icons allows users to infer the meaning of an object from its visual appearance.

Hutchins et al. (1986) go on to say that the success of DMI is related to the feeling of direct engagement that they produce in the user. Thus, the person feels as if she is interacting with the concepts of the domain rather than with an electronic intermediary. The shorter the semantic and articulatory distances, the greater the feeling of direct engagement. Therefore, the difference between DMI and conventional interfaces is that the gulfs are bridged by the computer, leaving the person to concentrate on the task at hand rather than the interaction process per se.

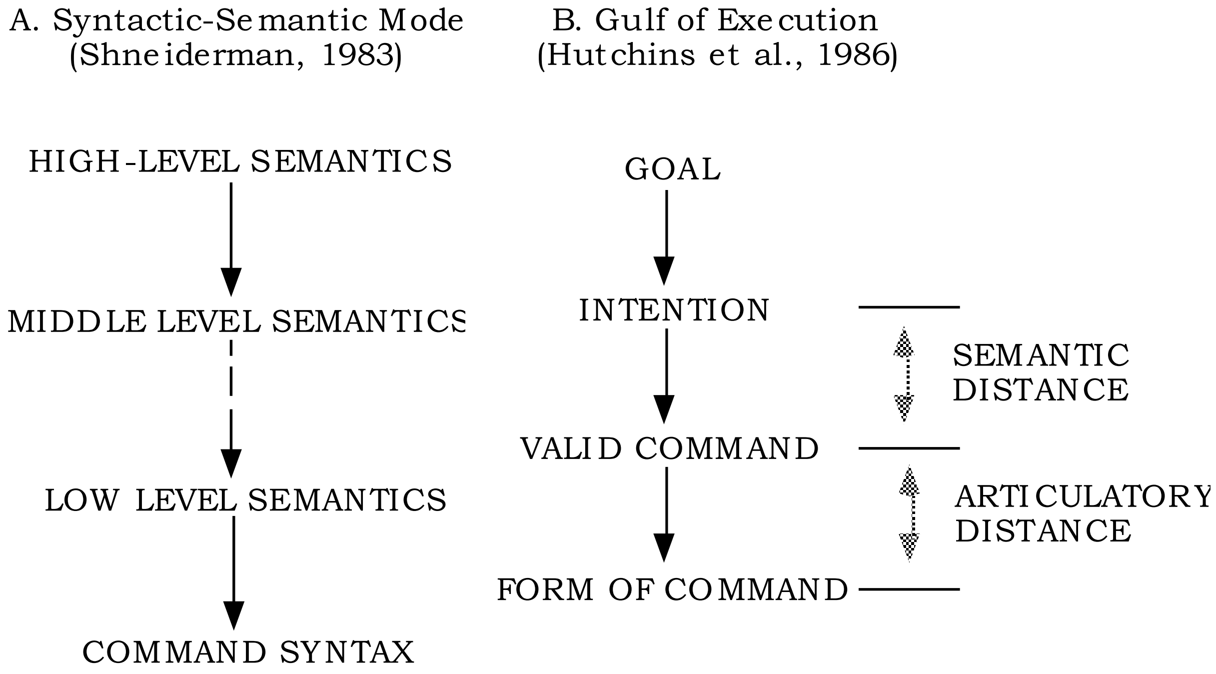

Figure 2. Process description of command formulation.

2.3. A Comparison of the Models

Obviously, the theory of Hutchins et al. (1986) is more elaborate than Shneiderman's (1983). Several of the limitations of the SSM that were identified above have been overcome. Firstly, the important distinction between perception (Gulf of Evaluation) and action (Gulf of Execution) is made. Whereas Shneiderman dealt almost exclusively with the formulation of commands, i.e., the semantic distance associated with the Gulf of Execution, Hutchins et al. provide a fuller account of the interaction process. A further improvement is the distinction between goals, intentions for accomplishing goals, and the evaluation of goals. Figure 2 shows the resulting contrast between the two theories in terms of their process descriptions for command formulation. Finally, the distinction is also made between the mapping between an intention and the meaning of an expression, and that between the meaning of an expression and its form. The end result is that the theory based on the Gulfs of Execution and Evaluation has greater psychologically validity than the SSM. However, as we shall see, even this account has its limits.

2.4. Limitations of Existing Theories With Respect to Process Control

One of the problems with the SSM is that it does not specify how the semantics should map onto the syntax. Hutchins et al. (1986) state that the main point to consider is the tradeoff between designing a command language that has generality and one that provides commands that are at a high level, closer to the way people think about the task. In the realm of process control, this is not a problem for interfaces are designed with very specific applications in mind; generality is not an important consideration. However, there are other properties of process environments that complicate matters considerably.

We will begin with the issues associated with the Gulf of Execution. Ideally, the mapping would map each intention onto a single unique action; this is the isomorphic mapping case. The only problem would then be to choose an appropriate level of abstraction, according to the tradeoff described above. However, an isomorphic mapping between intentions and actions is not feasible in process control because operators have a hierarchy of goals. Thus, the question arises whether there should be a command for each goal, regardless of its level in the hierarchy, or whether all the commands should be at some 'baseline' level. Also, there may be a one-to-many mapping between intentions and actions as a natural consequence of the complexity inherent in the problem domain (eg., coupling between subsystems). This is the well known problem of dealing with degrees of freedom. These complexities introduced by process control systems are beyond the explanatory capability of the model of Hutchins et al.

There are additional problems when we consider the Gulf of Evaluation. Figure 3 shows the different phases in the gulf. As with the Gulf of Execution discussed above, there are two segments: semantic and articulatory. Semantic distance is reduced by providing users with a powerful and veridical way of thinking about the problem domain. The display should be a representation of the domain. Hutchins et al. (1986) make the important point that multiple representations of the same objects are required to suit different goals, but they do not pursue the point to any great extent. But multiple representations are essential in process control, since as discussed previously, operators reason about the process at multiple levels. This would suggest that different representations may be required to support problem solving at different levels of abstraction.

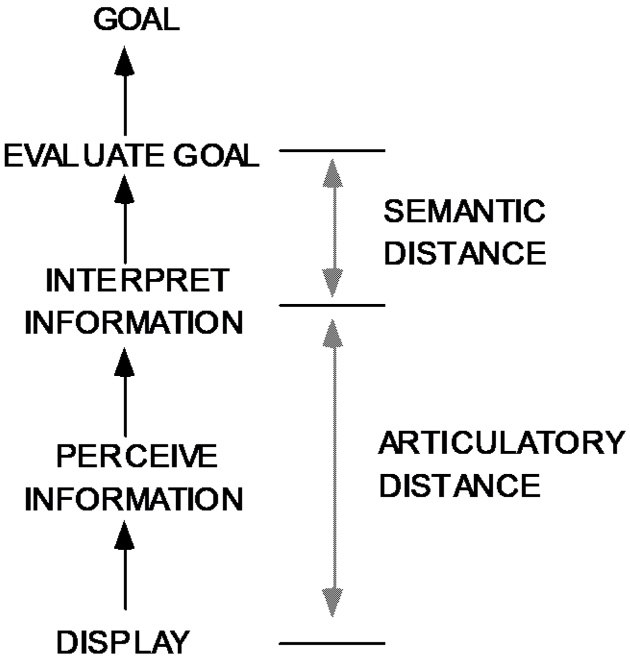

Figure 3. The different phases of the gulf of evaluation. Adapted from Norman (1986).

But even if we are dealing with a single representation, a given stimulus may be interpreted in different ways, depending upon the context, the operator's current goals, and his level of competence. Thus, the same display can be perceived as a pattern of signals for control of direct sensori-motor manipulations, as signs which may act as cues for release of heuristics, or as symbols for use in functional inference. These three semiotic categories provide the information necessary for control at the skill-, rule-, and knowledge-based levels, respectively (Rasmussen, 1983). During task performance, operators will process information at all three levels. But because each of the three levels of cognitive control generally supports different types of activities (this will be discussed in more detail in a later section), different information requirements are necessary. To be effective, an interface must be designed to support skill-, rule-, and knowledge-based processing. As mentioned above, the framework set out by Hutchins et al. (1986) does not take into account the concept of multiple levels of cognitive control.

In discussing the articulatory distance associated with the Gulf of Evaluation, Hutchins et al. (1986) limit themselves to the topic of icons. The question of form, however, is much more involved than merely dealing with the problem of getting the user to infer the intended meaning of an object from its form. Because we are mainly interested in communication via visual displays, we will discuss some additional considerations pertaining to the visual form of the displayed information.

Often, operators are required to make quantitative readings from a display in order to determine the state of a certain parameter. Thus, it is important to know what form the displayed information should take so that it is easy for the operator to accurately perceive the data. When multiple representations of the domain are implemented, the visual momentum supported by the system becomes an important consideration (Woods, 1985). The problem is to develop information presentation techniques that will reduce the mental effort required of users to integrate information across successive displays. Finally, the problem of perceived structure is also relevant. Any display has an inherent perceptual structure imposed on it by the human visual processing system. Thus, it would be useful to design the display in such a way that the perceived structure matches the organization inherent in the problem domain.

2.5. Summary

To summarize, existing theories of DMI have certain limitations with respect to process control. The fact that people can have a hierarchy of goals, and therefore reason at various levels of abstraction, is not accounted for in either of the theories reviewed above. The concept of multiple levels of cognitive control is a further complication. In addition, the issues associated with visual form have not been thoroughly addressed. What is needed is a a theory that takes into account these complications so that the benefits of DMI can be transferred to process control. The remainder of this paper is concerned with the development of such a theoretical framework.

3. A MODEL OF HUMAN INFORMATION PROCESSING

Our approach will be to develop a framework for interface design that supports the basic properties of the human cognitive system. By basing design decisions on knowledge about human behavior at the process level, rather than the performance level, the recommendations that are obtained should generalize across different applications (Pew and Baron, 1983). Thus, the first step is to describe what we know about human cognition.

We have adopted Rasmussen's (1974; 1986) model of human information processing (HIP) for this purpose. There is an important reason for choosing this particular model rather than any of the other existing models of HIP. Rasmussen's model was developed by observing operators' problem solving activities in actual work situations, and by analyzing human error reports from industrial accidents. Since we are interested in designing interfaces for complex systems, it is important that the model we select be grounded in the type of real-world situations that we are attempting to design for.

Only a brief overview of the model will be provided here. A more in depth description can be found in Rasmussen (1974; 1986). The model consists of two different subsystems (see Figure 4): a conscious processor and a subconscious processor, corresponding with attentional and automatic control modes, respectively. Such a dual architecture is common to most well accepted models of human performance (Reason, 1988).

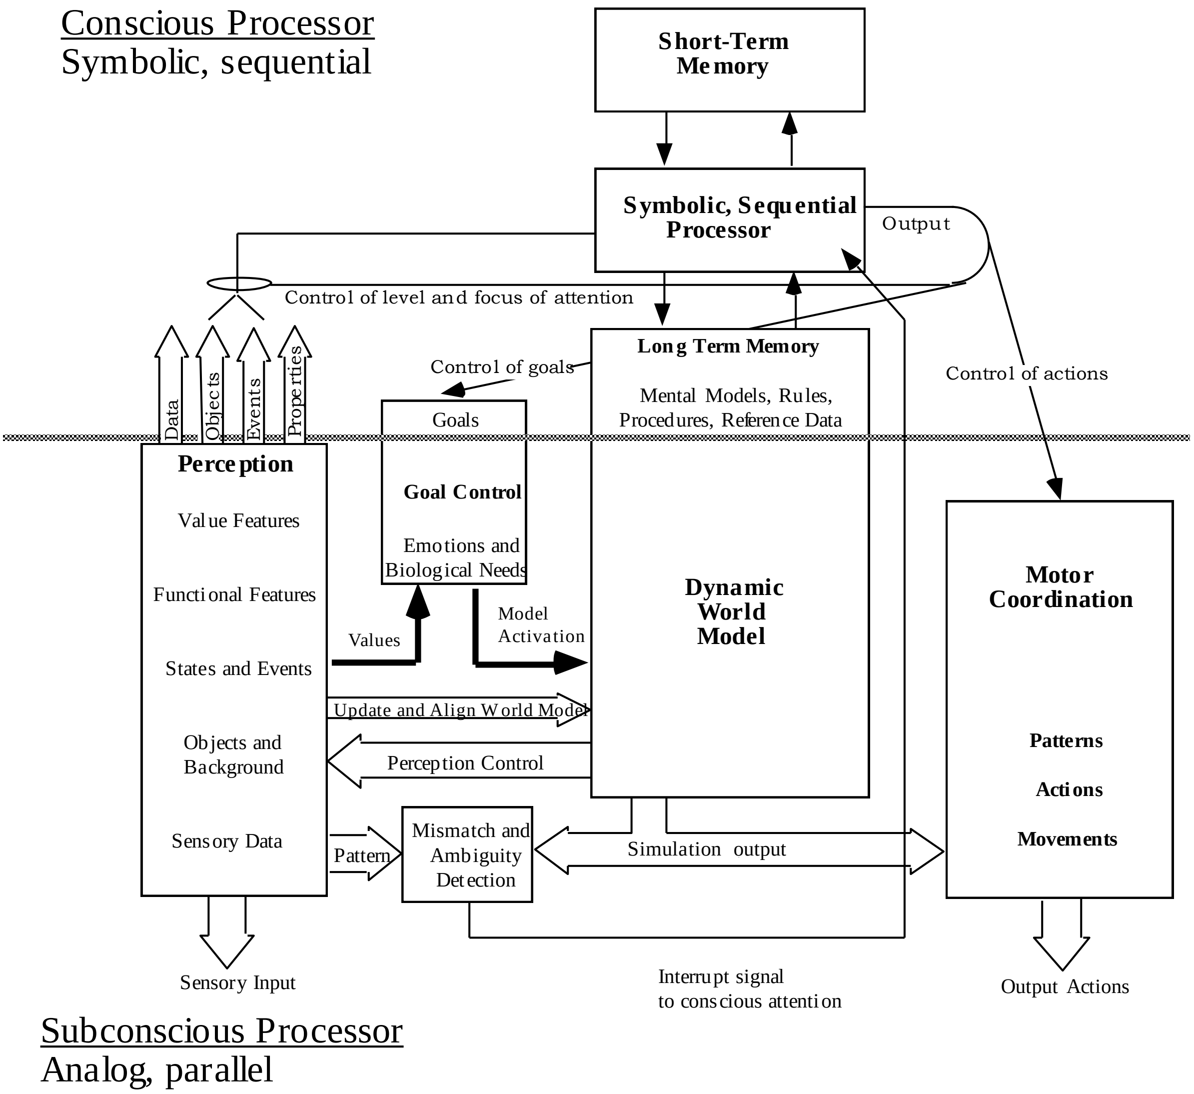

3.1. The Conscious Processor

The two processors have very different properties. The conscious processor is slow, sequential, effortful, and limited by the capacity of short-term memory. However, its strength lies in its ability to operate on symbolic information using its large repertoire of data-processing models and strategies. Such an architecture provides the flexibility that is required to cope with novel situations. The conscious processor is responsible for what is commonly referred to as rational or analytical thinking, i.e., activities such as improvisation, decison making, and symbolic reasoning. In addition, it functions as a high-level coordinator of the subconscious subsystem, which is the main data processor.

3.2. The Subconscious Processor

The subconscious processor is fast, parallel, effortless, and has a very high capacity for processing data. Basically, this subsystem is responsible for the functions of perception, motor control, and intuitive judgement. Information is decoded from the sensory input, and higher-level features are extracted from these data. The patterns extracted from the input data are then used to control the actions of the motor system. The coordination between sensory input, motor actions, and the dynamic environment is accomplished by a dynamic world model, which operates in real time. An essential aspect of this model is that it portrays the person as an adaptive, goal-oriented being. Thus, the dynamic world model enables people to attend to the invariant features of the environment that are essential to the actual goal in the present context (the role of perception), and to dynamically generate feedforward patterns of control actions (the role of motor control). In order to carry out this sensorimotor function, the subconscious processor deals with data in terms of time-space signals. The major limitation of this subsystem is that it is only capable of dealing with familiar, frequently encountered situations. When novel events are encountered, there is a mismatch between the behavior of the environment and the expectations provided by the dynamic world model. These conditions force processing to the conscious, symbolic level.

3.3. A Comparison of the Two Subsystems

The distinction between the conscious and subconscious subsystems of human cognition described by Rasmussen (1974) has been independently described in the literature under a variety of labels: smart vs. rote instruments (Runeson, 1977); structure oriented vs. state oriented response (Rouse, 1983); abstract reasoning vs. concrete activity (Chapman and Agre, 1986); analytical vs. intuitive cognition (Hammond, Hamm, Grassia, and Pearson, 1987); discrete symbolic mode vs. continuous dynamic mode (Carello, Turvey, Kugler, and Shaw, 1984); and recognitional vs. analytical decision making (Klein, in press). While differing in the particular details, all of these dichotomies are generally similar to the distinction described in Rasmussen's model. In this section, we are interested in identifying the relative merits of these two processing modes.

In effect, there is a tradeoff between subconscious and conscious processing. The subconscious processor provides a very efficient means of processing information, given that we are familiar with the characteristics of the environment, as represented in the dynamic world model. This type of processing is not possible for novel situations since the requisite model of the environment will not be available. The conscious processor, on the other hand, allows us to cope with new and unexpected situations but in a comparatively laborious fashion. The cost of being able to reason at an abstract level is that processing is slow and effortful. Thus, the two subsystems are complementary: neither one has a global superiority over the other, and both are necessary to cope with the entire range of demands that the person is likely to encounter. Together, they constitute a cognitive architecture that can tradeoff processing efficiency for the ability to deal with unanticipated variability.

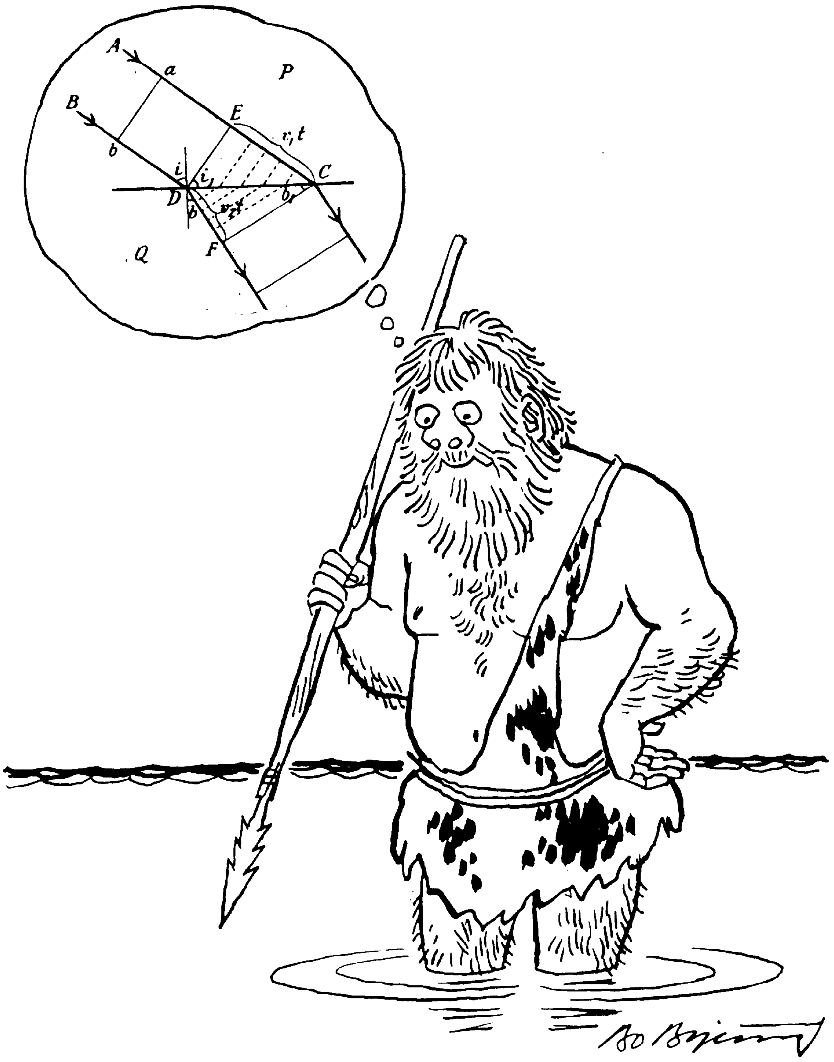

Figure 5. The cave man. Drawing by Bo Bojesen. Reproduced by permission.

The relative merits of these two processing modes is illustrated in Figure 5. A caveman is faced with the task of catching fish with a spear, and he has two different ways to go about it. If he were to rely on the analytical approach based on trigonometry, as shown in the diagram in the bubble, his chances of catching a fish would be negligible. The dynamic demands of the situation far exceed the processing capabilities of the conscious processor. However, with sufficient practice, the caveman will be able to acquire the skill of catching the fish by perceiving and acting, rather than thinking. His experience will allow him to develop a dynamic world model of the relevant factors: the optical physics, motor control of the spear, and movement of the fish. Once such a model is developed, the caveman will be able to effectively catch fish by relying on the efficiency of the subconscious processor. Clearly, this is a much more appropriate way of dealing with the demands of the situation.

However, as mentioned above, there are some disadvantages to the subconscious mode. One important point is that of experience. The caveman will have to have a considerable amount of practice with the task before he develops the necessary skills that will allow him to effectively rely on subconscious processing. Secondly, if for some reason, the demands of the situation change in a substantial way, the skills that the caveman has developed will no longer be appropriate. As an example, consider what would happen if the refraction index of water were to change significantly. The caveman's dynamic world model would no longer be appropriate and he would not be able to catch any fish. This is where the advantage of conscious processing comes in. Using the analytical approach based on trigonometry, a change in the refraction index would be easily accommodated merely by substituting a different value for that parameter in the equation.

3.4. Implications for Interface Design

What implications does this dual architecture model of HIP have for interface design? There are two characteristics of complex system domains (e.g., process control) which allow us to extract some interesting design implications from the cognitive model described above. First of all, the operators of such systems are highly skilled and have extensive experience in controlling the system. Secondly, interface design for complex systems consists of specifying an interface for a single, specific application; generality is not important. Thus, unlike word processors which will be used for a variety of applications, complex systems are designed with fixed and well defined purposes in mind. In this way, the problem domain can be thought of as a closed world.

These two factors, skilled operator and bounded problem solving world, make subconscious processing an attractive possibility. The fact that operators will have extensive experience with the system means that they will have the opportunity to develop a dynamic world model that will allow them to effectively rely on subconscious processing. Also, the fact that the operators will be dealing with the same system means that, if a way can be found to describe the problem solving world, then the need for dealing with novel situations should be minimal. This, in turn, implies that the reliance on conscious processing should also be minimized. Together, these two characteristics of complex systems suggest that the design should be aimed at taking advantage of the processing efficiency of the subconscious processor.

The characteristics of the subconscious processor have evolved from primitive man's necessity to dynamically interact with his environment in order to survive. These demands have resulted in the development of a subsystem that is characterized by very efficient feature extraction and classification, and dynamic coordination of the motor system with the environment. These are the skills that people find the most natural, and consequently, they peform them subconsciously, effortlessly, and in parallel. This is not to say that subconscious processing is perfect. Even in cases where people have developed the necessary skills, subconscious processing is still susceptible to errors, especially failures of attention (Reason, 1984). However, under certain conditions, the efficiency of subconscious far exceeds that of conscious processing. The claim that subconscious processing can be carried out more effectively and with less effort than conscious processing is a well founded one; it has both theoretical (Rasmussen, 1986; Reason, 1987; 1988) and empirical support (Hammond et al., 1987). An obvious goal for interface design would be to take advantage of this strength (cf. Brehmer, 1986).

It is important to realize that we are <u>not</u> claiming that, in situations where skilled people are interacting with a closed and fixed problem solving world, subconscious processing is always superior to conscious processing. As we will discuss below, this is certainly not the case. What we are claiming is that these conditions are propicious for subconscious processing and that the designer should take this into account in specifying the interface for the system. Given the proper interface, it should be possible for an experienced operator to effectively meet most of the task demands by relying on subconscious processing. The hypothesis behind this claim is that the level of performance of the man-machine system will be enhanced if the interface is designed in a way that will allow the operator to take advantage of the efficiency of subconscious processing.

We will close this section with a final comment on the nature of the recommendation we are making. Again, figure 5 will serve as a useful example. In designing interfaces, we cannot change the properties of the domain nor the properties of the person. Both of these are fixed. For instance, in figure 5, the caveman must take into account the physical principles illustrated in the diagram inside the bubble. These physical principles represent the demands of the task, and if the fish is to be caught then they <u>must</u> be taken into account. This cannot be avoided. The degrees of freedom in design are in how to deal with these demands. The argument developed above can be viewed as an attempt at answering the question: What is the most efficient and most effective processing mechanism that people can use to cope with the demands of the domain? Answering this question allows us to formulate recommendations for designing interfaces that attempt to take advantage of the people's most powerful capabilities.

4. THE PROPENSITY FOR SUBCONSCIOUS PROCESSING

The discussion so far suggests that it would be highly adaptive to take advantage of the efficiency of subconscious processing. Interestingly enough, people naturally adopt such a strategy. There is ample evidence from psychological research to show that people strongly favor the utilization of the effortless, parallel, routines available to subconscious processing (Hollnagel, 1981; Klein, in press; Rasmussen, 1974; Reason, 1987; Rouse, 1983). People attempt to simplify complex tasks by taking advantage of their most powerful cognitive resource.

4.1. An Example

The work of Klein (in press) provides an excellent example of people's propensity for subconscious processing. He conducted a series of naturalistic studies of expert decision making in several domains: firefighting, military operations, and engineering design. The data were collected by, first, identifying non-routine events requiring skilled decision making, and second, conducting interviews to probe these events in order to examine the nature of the decison making process. Over a hundred cases were analyzed in total.

Since the incidents examined were non-routine, we would expect that decision making would be, to use Klein's terms, analytical rather than recognitional, i.e., based on conscious processing rather than subconscious processing. Surprisingly, the results indicated that, even in such critical incidents, experts relied mainly on the recognitional mode of decision making. Such a strategy is adaptive in several respects. In terms of mental effort, the recognitional mode is less taxing than the analytical mode. In terms of effectiveness, recognitional decision making allows experts to take advantage of their experience. Because of his wealth of experience, the expert is able to decide on an appropriate action to take rather than consciously generating and evaluating a set of alternatives. Finally, in terms of appropriateness, recognitional decision making is much quicker than analytical decision making, and therefore allows experts to effectively cope with time stress.

These results support the argument outlined in the previous section. If we can design interfaces that allow people to effectively take advantage of subconscious processing, then the benefits can be great. However, there is a major difference between the domains that Klein investigated and those with which we are concerned: in complex, high technology systems, the system being controlled is not directly observable. While Klein's results illustrate the potential benefits of subconscious processing, they do not give us any indication as to how to derive these benefits through proper interface design, nor do they tell us what can go wrong if the proper support for subconscious processing is not provided.

4.2. What Can Go Wrong

The problem with existing interfaces is that, rather than supporting this preference for subconscious processing, they penalize people for it instead. An experimental study conducted by Hollnagel (1981) in the area of process control provides an excellent example of this point.

The fact that the process being controlled is not directly observable leads to two phenomenologically different types of control strategies that can be adopted by operators. Following Hollnagel (1981), we will refer to these as surface control (corresponding to subconscious processing) and deep control of the system (corresponding to conscious processing). Surface control is guided by the plant as it is represented by the displays. In contrast, deep control of the system is guided by one's knowledge of the underlying process. While this surface control/deep control distinction is best thought of as a continuum, and not as two discrete control strategies, experience with investigations of process environments has often indicated that operators have a distinct preference for surface control rather than deep control of the system (Rasmussen, 1974; Rouse, 1983). This is how the omnipresent characteristic of human cognition to prefer subconscious to conscious processing manifests itself in process control.

Hollnagel's (1981) experiment provides a typical example of this pattern of behavior. Subjects tended to disregard the properties of the physical process which was controlled, and relied on the perceptual characteristics of the display instead. In effect, they often treated the system as if it was physically structured as the display indicated. However, as is the case with most existing systems, the displays were not designed to be veridical representations of the process. Thus, the strategy of surface control results in several classes of problems. First, it is easy for operators to forget, and therefore fail to consider, properties of the process which were not shown in the displayed representation of the system. Secondly, because of the inconsistent mapping between the abstract properties of the process and the cues or signs provided by the display, the cues that operators normally rely on to control the system are imperfectly correlated with the state of the system. Thus, in novel situations, surface control will result in underspecification and human error (Rasmussen and Vicente, 1987).

It is evident that present systems are not designed to support the basic characteristics of human cognition, as outlined above. This results from the fact that most current displays present the operator with elemental, physical data that are measured by sensors. Thus, in order to adopt a deep control strategy, the operator must perform what we call a translation task, i.e., he must map the elemental data that are displayed onto the abstract, higher-level properties of the process. Because this mapping is usually quite complex, the translation task requires considerable effort. In a similar vein, it is difficult for operators to map their intentions onto a sequence of actions that the system will accept. The important point to realize is that, due to the way systems are currently designed, deep control requires the operator to perform the translation task unaided. In keeping with the preference for subconscious processing, operators will often adopt a surface control strategy rather than take the effort to perform the translation task necessary for deep control.

4.3. Implications for Interface Design

In the previous section, we suggested that interfaces should be designed in such a way as to allow people to effectively rely on subconscious processing since it is more effective and less effortful than conscious processing. In this section, we have identified another reason for supporting subconscious processing: not only is it a more effective way of dealing with the demands of the task, but it is also the strategy that people naturally adopt. The work of Klein (in press) shows how prevalent this strategy is, as well as the benefits that can result from it. The work of Hollnagel (1981) illustrates the types of problems that are encountered if the interface does not support this strategy.

It is worthwhile mentioning that the idea of supporting operators' natural tendency to take advantage of the processing efficiency of the subconscious processor is not a novel suggestion, as the following quote from Rasmussen (1974, p. 11) indicates:

> "the process operator lives in a complex world. The information presented to him is a code for the physical, dynamical process in the interior of the plant. He is able to ...operate on the physical meaning of the symbols by rational deductive reasoning. During frequent routine tasks, however, he may be operating the control desk - not the process - by his subconscious routines. He may be able to improvise rapidly and subconsciously ... if he is allowed to break down subconsciously the information patterns into familiar generic units. This is only possible if he can control the process directly - the display system therefore should be 'transparent' and the physical process should be directly 'touchable' on the control desk."

The similarity between these ideas and the basic concepts behind DMI is striking.

We are now in a position to reformulate the claims for DMI in terms of the general properties of human cognition. The success of DMI can be attributed to the fact that they allow people to take advantage of the efficiency of subconscious processing. This is a theoretically satisfying explanation of DMI but it does not give us any indication as to how to go about building such interfaces. This is a result of the fact that the description of human cognition that we have used up to this point, while useful, is overly simplified. In order to develop a more complete understanding of the cognitive processes relevant to interface design, we require a more comprehensive framework. Therefore, in the nect section we will leave behind the conscious vs. subconscious distinction, and will instead adopt Rasmussen's (1983; 1986) skills, rules, knowledge (SRK) framework.

5. MULTIPLE LEVELS OF COGNITIVE CONTROL

5.1. The SRK Framework

The basic tenet of the SRK framework is that information can be interpreted in three distinct ways: as signals, signs, and symbols. The way in which information is interpreted determines which of the three types of cognitive control is activated: skill-based behavior (SBB), rule-based behavior (RBB), and knowledge-based behavior (KBB), respectively. Thus, control may depend on a repertoire of automated behavioral patterns (SBB), a set of state-action production rules (RBB), or problem solving operations in a symbolic representation (KBB).

In terms of the model of HIP outlined earlier, SBB and RBB are within the realm of subconscious processing whereas KBB is the responsibility of the conscious processor. We can reformulate the points made in the previous sections in terms of the SRK framework. First, lower levels of cognitive control tend to be carried out more effectively and with less effort than higher levels. Second, people have a definite preference for relying on lower levels of cognitive control. The basic implication that we derived is that interfaces should be designed to allow people to meet the demands of the task by relying on lower levels of cognitive control.

How does one go about doing this? In very general terms, we know that the form of the information must meet certain requirements for each level. SBB can only be activated when information is presented in the form of time-space signals. RBB, on the other hand, is triggered by familiar perceptual structures. And finally, KBB is activated by meaningful relational structures.

However, the form in which information is presented does not directly determine which level of cognitive control will be activated. Several other factors are also important. First, all other things being equal, more demanding tasks will tend to require higher levels of cognitive control (cf. Rasmussen, 1983; Sanderson and Harwood, 1988). Second, all other things being equal, the degree to which an actor can effectively rely on lower levels of cognitive control is also a function of his skill and experience. More experienced actors are able to deal with most task demands by relying on lower levels of cognitive control (cf. Dreyfus and Dreyfus, 1986; Klein, in press; Rasmussen, 1983; Sanderson and Harwood, 1988). In summary, the demands of the task, the actor's experience, and the form in which information is presented combine to determine which level of cognitive control is activated.

This means that even if we follow the recommendation made earlier of designing interfaces that allow people to rely on lower levels of cognitive control, they may nevertheless have to resort to higher levels of cognitive control. Therefore, merely supporting the lower levels is not sufficient. To be truly effective, an interface should also support higher levels of cognitive control. In order to determine what the information requirements for each of the levels are, it is necessary to understand how each of the different levels are related, and what the activities associated with each level are.

5.2. Interaction Between Levels

While it is possible to describe each of the levels independently, task performance will usually require a simultaneous consideration of all three levels of cognitive control (see Figure 6). For instance, during execution of skill-based routines (shown as synchronous activities in Figure 6), conscious attention is free to, and usually does, cope with other matters on a time sharing basis. As an example, if the next task requires a sequence of activities which will not integrate into an automated pattern, the rule-based domain will be involved in retrieving a relevant rule-set from memory as the skilled movements are being performed. Also, during task performance interrupts may occur when choices are to be made or when adjustments of the current internal model are needed. These rule-based behaviors are shown in Figure 6 as synchronic activities.

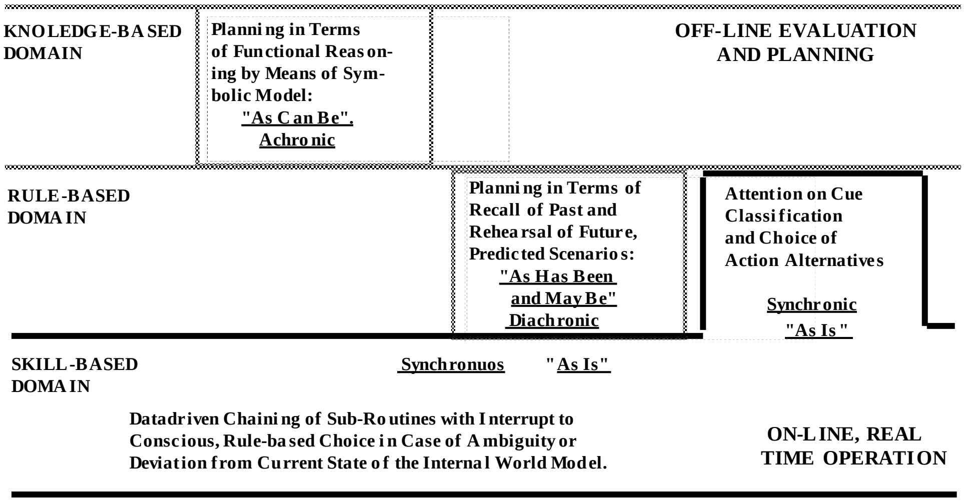

Figure 6. Interaction between the different levels of cognitive control.

When skilled activities are running smoothly, attention can also be diverted to the evaluation of past and planning of future activities. This may require knowledge-based analysis and planning, or recall and evaluation of the success of previous application of rules. These are shown in Figure 6 as achronic and diachronic activities, respectively. Important activities include forecast of future needs, recall and internal rehearsal of relevant rules, and arrangement of rules in proper order in memory for later activation. For the present purposes, the most important implication of this dynamic interaction between the different levels of cognitive control is that the information required for on-line control of the current activity, and off-line planning considerations may not belong to the same time frame, nor to the same part of the problem space.

It is evident then that the structure of the internal control of cognitive activities has important implications for interface design. The information presented to the operator will have at least three distinct functions in the control of a complex work sequence. Information must be available for activation of skilled routines; for control of the course of the routines and, finally, for monitoring of proper result of an activity. Each of these will be discussed in greater detail below. To anticipate, the important point is that it is essential that an interface be able to support control of performance at different cognitive levels, not only because users with different levels of training may use the same system, but also because activation, control, and monitoring will require different interactions between cognitive levels.

5.3. Activation of Task Routines

Routines will typically be activated by stereotypical signs that are empirically correlated with a call for the various task routines. Signs are selected from the available information to be clearly salient features which are convenient and sufficient for discrimination between the usually relevant alternatives for action. It is a common experience that such signs very often are informal information, such as noises from machinery and the sign-action mapping will typically only be valid under 'usual work conditions' leading system users into traps during less familiar situations. An advantage offered by advanced information technology is that signs that are useful for activating sub-routines can be designed as patterns with a more reliable mapping onto task requirements.

5.4. Control of Activities

Proper control of movements during skill-based routines depends on proper alignment and synchronization of the dynamic world model, a function requiring perceptual information patterns serving to maintain the quantitative time-space signal transmission. A match with the active mental representations at this level requires animated, quantitative, and graphical presentation of 'relational structures' (Craik, 1943).

At the rule-based level, we are concerned with supporting memory for sign-rule correlations. It is important to take into account the fact that infrequent but risky conditions may require a certain rule configuration which may not be communicated by some sign configurations. This leads to errors due to 'procedural traps' (Rasmussen, 1987). A proper match at this level in terms of information support will be against a kind of mental decision table. It will be important to support the evolution of cues which are not only conveniently perceptive, but also reliable across work situations.

At the knowledge-based level, control of the course of reasoning will require presentation of symbolic information with reference to the mental model of the 'relational structures' of the work content, e.g., to maps of system anatomy suited for causal reasoning. A fundamental problem is that the relational structure of a task content can be represented at several levels of the means-end dimension. Preparing suitable control strategies for unusual work situations is a design task, and depends on the ability to operate in a means-end/part-whole problem space. In complex systems, multiple display representations may be needed to represent the problem at each of the various levels of abstraction.

5.5. Monitoring of Performance

For effective error recovery, it is important that performance monitoring be performed without undue delay. The problem is that, typically, work performance depends on a long sequence of actions. Thus, monitoring of performance with reference to the ultimate goal will be inefficient because the feedback will probably be received so late that the process has reached an irreversible state. Consequently, monitoring must depend on some kind of evaluation of the process of attaining the desired goal. The reference for judgement will depend on the circumstances. For instance, during skill- and rule-based control of familiar tasks, acceptable performance can be judged with reference to the usual, normal response of the environment, i.e., on monitoring normal information feedback patterns. This means, that information defining normal responses to interaction should be coded into integrated perceptive patterns that are easily recognized, that are easy to detect changes in, and above all, that effectively discriminate between different work conditions.

During less familiar rule-based activities when perceptive references in terms of normal patterns are not available, monitoring depends on concurrent knowledge-based analysis of the response during the task sequence. Thus, the operator must 'understand' the responses of the task environment by referring to his mental model of the process. Typically, knowledge-based behavior involves planning and control of a sequence of actions to reach a chosen goal or product. As mentioned above, these activities will require mental models quite different from those necessary for control of the proper course of actions during the process.

5.6. Summary

Let us summarize the argument developed in this section. First, we have argued that performance of real world tasks requires a complex interaction between the different levels of cognitive control. We also found that the interface alone does not determine which level of cognitive control will be triggered. The experience of the person and the demands of the task are also contributing factors. Together these two points indicate that all three levels of cognitive control will be invoked during task performance. In addition, we have demonstrated that each level requires a different type of information support, and that KBB, because it is so effortful, requires appropriate support if it is to be carried out effectively. The primary conclusion that emerges from this discussion is that an interface should support all three levels of cognitive control.

Based on the theoretical development conducted to this point, we can formulate a specific goal for interface design:

An interface should not force cognitive control to a level higher than the demands of the task require, but at the same time, it should provide the appropriate support for all three levels of cognitive control.

The obvious question, which will be addressed in the following section, is: How does one go about designing an interface that will accomplish these goals?

6. A TAXONOMY OF HUMAN-SYSTEM INTERACTION MODES

In this section, we present a morphological analysis (cf. Zwicky, 1967) of the transformations necessary for matching the interface representations to the requirements of each of the three levels of cognitive control.

In our taxonomy, we distinguish between four different types of communication: direct, signal mapping, sign transformation, and symbolic transformation. Each of these can be used to describe the transformations in the sensory path (perception) or in the motor path (action). A factorial combination of all possible cases results in sixteen categories of human-system interaction, as shown in Table 1. The modes with symbolic action transformations are segmented from the rest of the table because, with present technology, it is impossible to control a system via symbolic acts. Only for intelligent systems, which understand the meaning of messages (i.e., that perform under the control of conceptual models and goals, rather than rules), can input acts be interpreted as symbols. Compare this with Searle's (1981) discussion on whether computer programs are able to understand language.

	Transformation Across Manipulation Interface			
Transformation across observation interface		2. Monatomic transformation of time-space properties of movements (signals).	3. Recoding of messages and movements into command (signs) modifying system functions.	4. Transmission of symbols to be interpreted by intelligent agent with goals and 'mental' model
1. Direct observation	1.1 Eating lunch; sculpturing; mowing objects; walking around.	1.2 Car driving; operating a bulldozer with analog joystick;	1.3 playing musical instrument with keys; operating crane with control keys, operating TV-set;	1.4
2. Data to be manipulated represented by analog space time configuration	2.1 eye surgery through microscope.	2.2 Operation of tele-robot through joy-stick and closed circuit TV; Driving tank through infrared goggles. Controlling process by direct manipulation of graphic analog display	2.3 Operating robot by command keys and TV. Controlling process from graphic analog display by command keys and switches.	2.4 This column is presently not relevant unless a human agent is a part of the system and interprets messages, i.e., transforms to column 1-3.
3. Data coded in terms of signs with reference to coding convention	3.1 Positioning object according to gestures; blindfolded walking from instructions.	3.2 Blind landing of aircraft from traffic control instructions.	3.3 Operating crane by control keys and signs (gestures).	3.4
	possible only when gestures and instructions are analog signals			
			Process operation by keys and indicator lights;	
4. Data coded symbolically with reference to conceptual model	4.1	4.2 Not applicable	4.3 Process operation, interpreting displays; operating switches and keys;	4.4

TABLE 1. A taxonomy of human-system interaction modes.

In this section, we will describe a few of the categories in Table 1, and provide some examples. It is important to note that any given interface may support several interaction categories. However, as mentioned above, an interface's effectiveness is determined by allowing the operator to exert lower levels of cognitive control during most of the time, and only require him to go to higher levels when the circumstances (i.e., the demands of the domain) warrant it. Thus, successive levels of interaction can be thought of as incremental in the sense that they subsume the lower levels. Finally, it should also be noted that the lower levels of interaction are concerned only with surface control (i.e., with the physical form of the environment), while higher forms of interaction are concerned with both surface and deep control (i.e., with the functions of the internal process being controlled as well).

6.1. Direct Observation - Direct Control (1-1)

This interaction mode relies on a closed time-space signal transmission path through the sensorimotor system. Behaviour deals with control of the form, location, and configuration of physical objects in the material environment. Typical examples are sculpturing clay by hand and eating one's lunch. Quantitative data on time-space features of the environment serve to align, update, and synchronize the internal world model. For instance, seeing that one's cup is filled to the brim will serve as a sign to change one's behavior (i.e., to switch from a 'normal drinking' internal model to a 'careful drinking' internal model). Also, features in the environment serve as signs, empirically correlated with the need for updating the internal model. Control is expressed at the level of intentions (e.g., be careful since the floor is wet), but the actual carrying out of the actions is subconscious and effortless. The definition of signs is performed through experience, and at this subconscious level, will be purely inductively correlational and probably not accessible to conscious report.

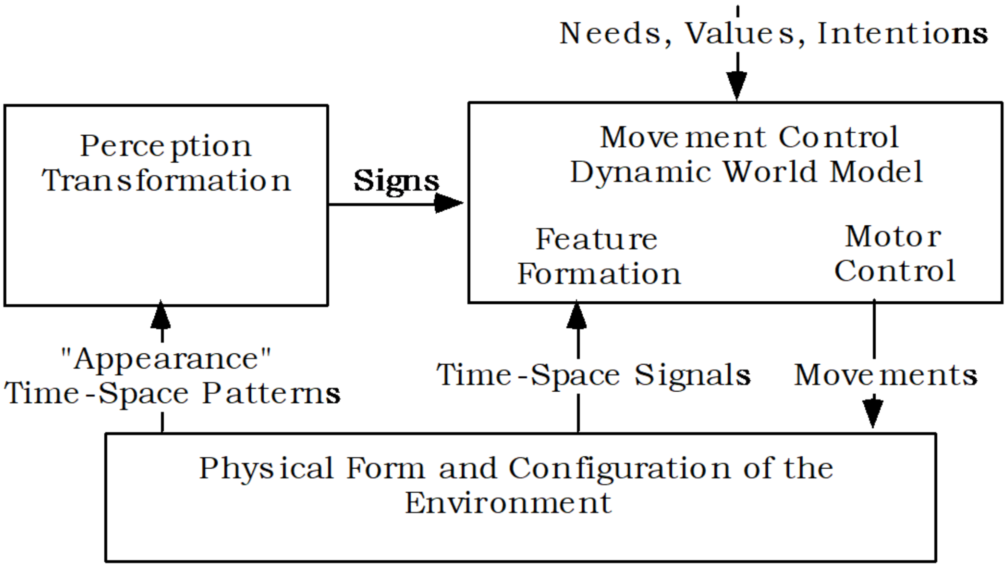

Figure 7. Direct observation and direct control: 1-1.

6.2. Direct Observation - Indirect Control (1-2, 1-3)

In both of these categories, a tool is inserted between the operator and the work context. The difference between the two cases is in the type of transformation that the tool provides.

1-2. The prerequisite for this mode of interaction is that the transformation in the motor path must allow for the transmission of actions in terms of time-space signals, thereby enabling subconscious (i.e., analogue) processing. Thus, the control structure from category 1-1 is unchanged because the time-space communication of quantitative signals through the tool is maintained. However, different kinds of transformations may be involved, such as the amplification of movements or of force applied, or a coordinate transformation (e.g., cartesian / polar). Examples of this mode of interaction are driving a car, or operating a bulldozer. In phenomenological terms, the tool is embodied as an extension of the motor system. Attention is not on the interface between the tool and the body, but rather, on the interface between the tool and the work context.

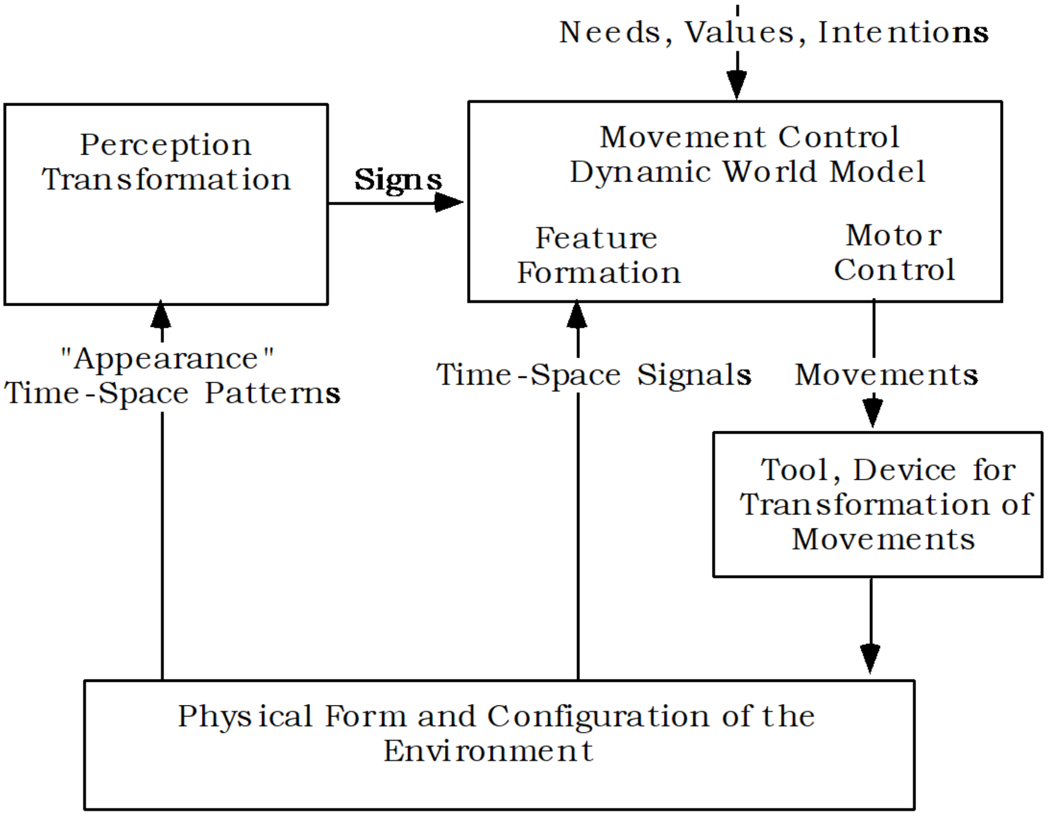

Figure 8. Direct Observation and Indirect Control, 1-2.

<u>1-3</u>. In this case, the tool does not accept commands in the form of time-space signals, but rather in the form of signs. The transformation from movement of the tool to action on the environment is a matter of pure convention. Examples are playing a musical instrument with keys, or operating a crane via a keyboard. With this category of interaction, the tool is no longer perceived as an extension of the body. Instead observation of the surface of the tool is also necessary because the operator must translate his intentions into a form that is acceptable by the tool's command language. As a result, attention must be divided between activities dealing with the intended effects of an action on the environment, and those dealing with the manipulation of the tool itself. The high capacity sensorimotor function will on-

ly serve the tool interface manipulation task (communicating commands, or signs, for change in system function). With highly experienced operators, however, subconscious processing may be possible if an intention can be carried out by an automated subroutine of control actions. If this is the case, then the translation task is no longer necessary, and behavior is similar to category 1-2. It should be noted that such skilled behavior can only be achieved with a great deal of practice, and even so, only for situations that are frequently encountered.

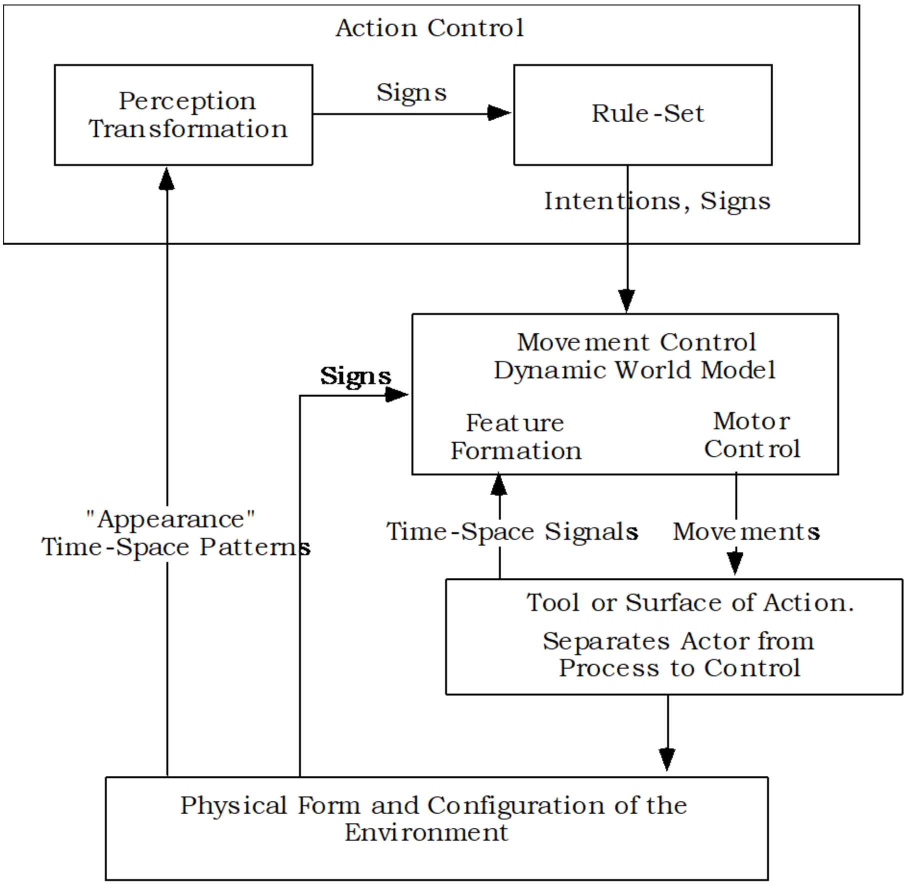

Figure 9. Direct Observation and Indirect Control: 1-3.

There are cases where it may be difficult to classify an interaction mode as either of type 1-2 or 1-3. In particular, transformations including either differentiation or integration (i.e., changing the order of control) may be characterized as a signal or a sign. For instance, if an on-off switch is used to control a movement by means of a motor, the switching movement may be defined as a sign, but through the integrating function of the motor, the time-space, quantitative loop may still be considered intact as long as there is a one-to-one mapping between movements and the effect on work content. An example of this situation would be control of a crane from a set of function keys.

6.3. Indirect Observation - Direct Control (2-1, 3-1)

2-1. In this case, a transformation device transmitting time-space signals is inserted in the sensory path. The observed information is an analogue representation of the environment in which acts are to be performed. Because the time-space loop is unbroken, the equipment is perceived to be an extension of the senses, thereby allowing subconscious processing to take place. As a result, attention is devoted to the actual state of the work context rather than on interpreting the displayed information. Examples of this interaction mode are eye surgery through a microscope, or walking at night with infra-red goggles on. The transformation in the sensory path need not be isomorphic; a homomorphic transformation will suffice as long as the critical time-space signals are communicated. An example would be the mirror symmetry between hand movements and visual experience when working with a microscope. Therefore, other transformations such as a change in frame of reference (e.g., cartesian, polar, etc.) may be included as long as objects, and time-space properties are consistently mapped.

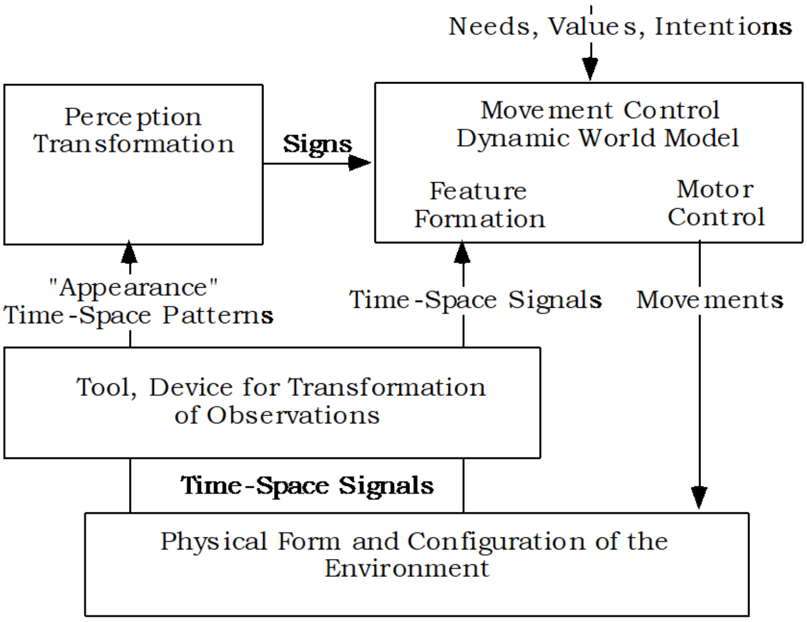

Figure 10. Indirect Observation and Direct Control: 2-1.

3-1. If the analogue representation is not maintained, and the mapping of time-space data and objects is interrupted, sensorimotor performance is no longer possible. Instead, the person must rely on rules for interpreting the displayed information. Examples of this interaction mode include a pilot being guided down to land via radio instructions, or cooking from a cookbook for the first time. In this case, the communication channel is no longer embodied as part of the senses. The high capacity perceptual functions can be used for perceiving the displayed information, but the operator must perform the translation task in order to determine what the arbitrary display conventions mean in terms of what is going on in the work context, or alternatively, adopt a surface control strategy and rely only on signs to control the system. Again, with experience the effort associated with the translation task can be reduced, but only for frequently encountered situations.

6.4. Indirect Observation - Indirect Control (4-3)

Examples of this type of interaction include control of 'invisible' processes like chemical process plant control. In these situations, two levels of control are relevant: the actions on the control panels (surface control), and the control of the process itself (deep control).

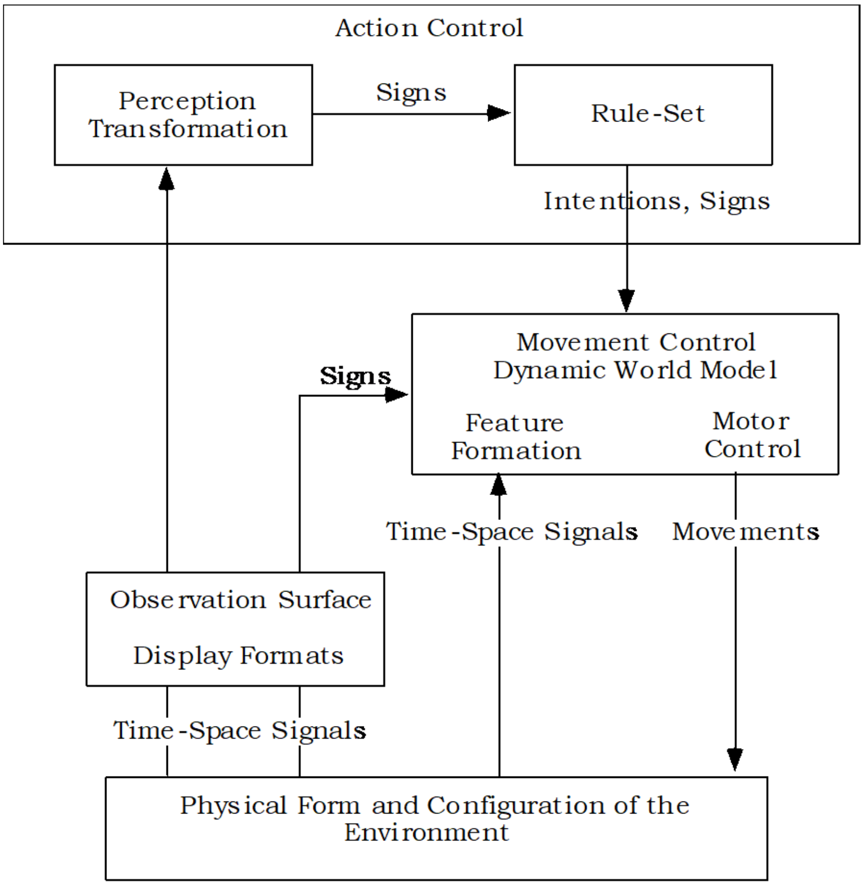

Figure 11. Indirect Observation and Direct Control: 3-1.

Typically in present systems, the observation surface and the action surface are physically separate (e.g., meters and screens as displays, and keyboards and switches as controls). On the action side, motor control is only concerned with operation of the keys and switches which send signs to the interior process by means of a one-way command language. On the observation side, displays present individual variables representing states of physical processes. Taken together, these characteristics of current systems force the operator to interpret the displayed information into meaningful higher-level concepts, and also to translate her intentions into action sequences that conform to the computer's command syntax. In other words, she must perform the translation task unaided.

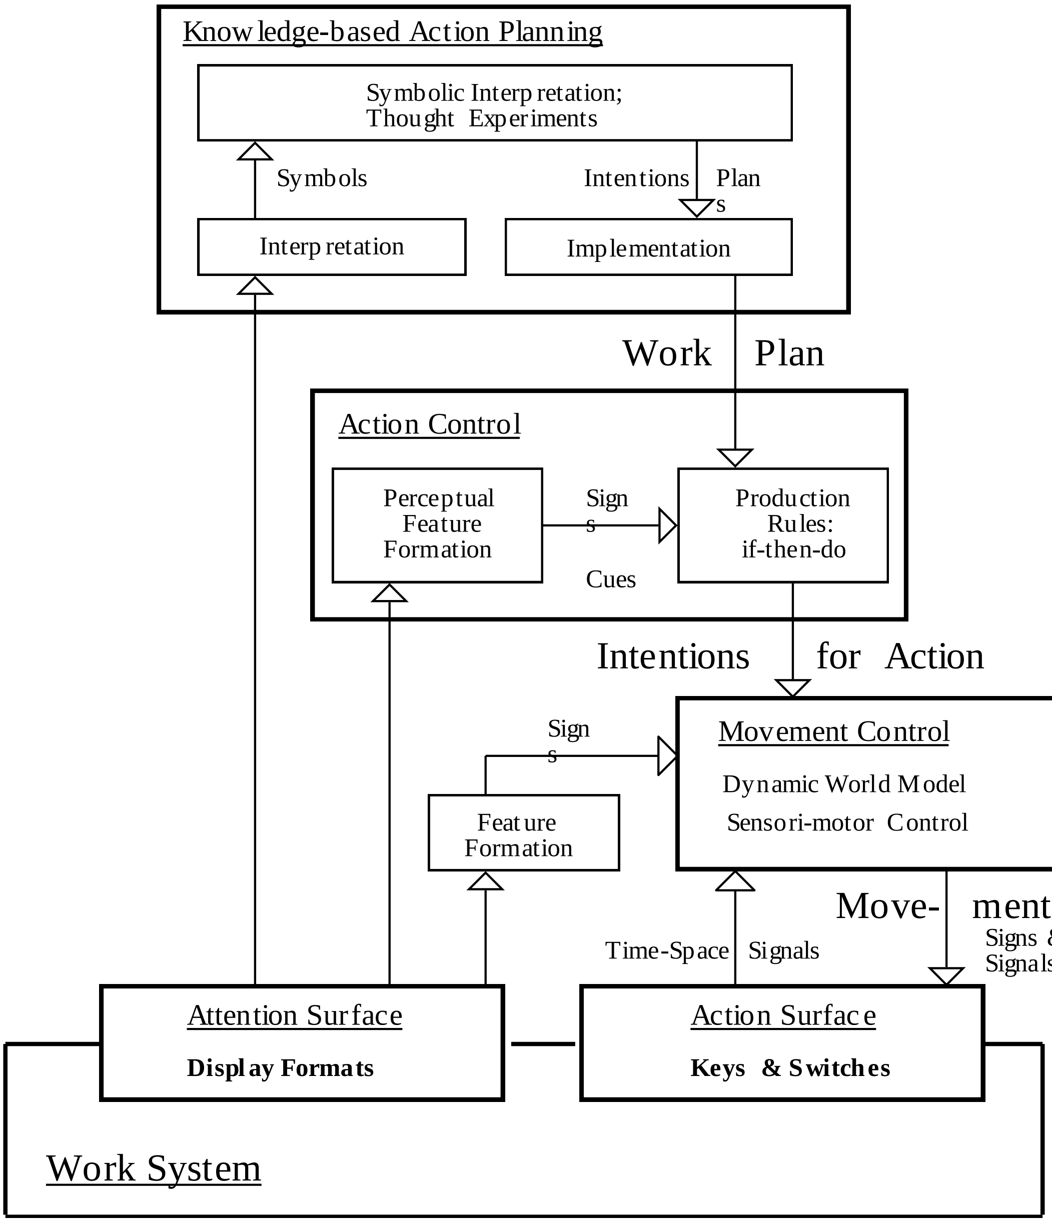

Figure 12. Indirect Observation and Indirect Control, 4-3.

The important point to realize is that it is possible to design an interface in such a way that the translation task is performed by the designer and the computer, rather than by the operator. In order to design such an 'ecological interface', the observation and action surfaces must be merged so that the time-space loop is maintained, thereby enabling subconscious processing. In addition, it is also necessary to develop a consistent one-to-one mapping between the abstract properties of the internal process to be controlled and the cues provided by the manipulation/observation surface. The goal is to make the invisible, abstract properties of the process (those that should be taken into account for deep control of the process) visible to the operator. In semiotic terms, this means that the cues provided by the interface have a consistent mapping onto the process properties. In this way, the same conceptual model may act as a symbolic representation when considered in relation to the elements of the environment and the laws controlling their relationships, and as a system of prescriptive signs when considered in relation to the rules for model transformation and data processing.

7. ECOLOGICAL INTERFACE DESIGN

In the previous section, we discussed the transformations necessary to match the interface representations to the different levels of cognitive control. Based on this analysis, we propose the following prescriptive design principles that constitute the EID framework:

1. Synthesize the control and the observation surfaces so that interaction can take place via time-space signals.

2. Have the computer perform the translation task by

developing a consistent, one-to-one mapping between the invisible, abstract properties of the process and the cues or signs provided by the interface.

3. Display the process' relational structures in the form of a functional hierarchy to serve as an externalised mental model that will support knowledge-based processing.

A major benefit of the EID approach which is not discussed in this paper is that, because it attempts to support all three levels of cognitive control, EID minimizes the likelihood of errors related to interference between internal control structures. This aspect of EID is described in more detail by Rasmussen and Vicente (1987).

7.1. An Example

In order to make these abstract principles more concrete, we will provide an example of an ecological interface. Figure 13 illustrates such an interface. A thermodynamic overview of a process system is presented in terms of a Pressure vs. Temperature graph, commonly known as a P-T diagram. The two dark lines plotted on the graph show the time history of the hot and cold legs of the process. The top end of the line represents the current state. The dotted lines indicate the critical limits for the variables. Thus, the operator's task is to keep the state of the two legs within the critical limits.

Following the first principle of EID, the operator is able to act on the interface itself. Thus, in order to change the state of one of the legs, the operator merely points to the end of the trend curve with a mouse and drags the cursor in the direction that he wants the system to go in. Another interesting feature of this interface is that the abstract properties describing the process' behavior (i.e., the thermodynamic relationship between temperature and pressure) are mapped onto the interface's perceptual characteristics. This means that the operator can perform the control task without having to resort to higher levels of cognitive control. Thus, the goal of the operator can be reformulated simply as: keep the endpoints of the two dark lines within the boundary defined by the dotted lines. In this way, a complex thermodynamic control task can be performed as a tracking task. The operator need not consider what the signs mean; the axes might just as well represent apples and oranges rather than pressure and temperature. On the other hand, if for some reason, cognitive control is forced to the knowledge-based level (e.g., if an abnormal situation is encountered, or if the operator is relatively inexperienced) then the very same interface can be interpreted as symbols for conscious reasoning. In this mode, the operator would consider the semantic content of the interface, (i.e., the thermodynamic relationship between pressure and temperature) rather than rely merely on its perceptual characteristics.

Several qualifications are in order. First, it should be noted that the figure is intended merely as an explanatory device and not as an actual design proposal. The interface has not been constructed nor evaluated. Secondly, the interface represents the process at a single level of abstraction. A complete design would necessarily include other levels of process representation.

To conclude, we see that the hypothetical interface shown in figure 13 embodies all of the principles of EID. It is consistent with the goal of presenting information in a way that allows people to take advantage of the efficiency of lower levels of cognitive control, while at the same, attempting to provide the necessary support for higher levels as well.

8. COGNITIVE CONTROL IN VARIOUS WORK DOMAINS

In this section, we provide an analysis of the cognitive control mechanisms associated with three domains: process control, musical skill, and bibliographic search. The purpose of these analyses is to show that the EID framework provides a fruitful perspective for analyzing work domains with an eye towards interface design.

8.1. Process Control

Figure 14 illustrates the mappings between the process, the interface, and the operator's mental model for a typical process system. The activities associated with each of the three levels of cognitive control are described below.

8.1.1. Skill-based level. Because the operator cannot directly observe or act on the process, the sensorimotor control patterns at the skill-based behavior level will only be concerned with the manipulation of items on the interface surface. The use of a mouse or a trackerball is preferred to command languages for this task because it maintains the communication of spatial-temporal aspects of the perception-action loop intact. To allow the development of a high degree of manual skill, the interface must be designed in such a way that the aggregation of elementary movements into more complex routines corresponds with a concurrent integration (i.e., chunking) of visual features into higher level cues for these routines. Thus, the display of information should be isomorphic to the part-whole structure of movements rather than being based on an abstract, combinatorial code like that of command languages.

8.1.2. Rule-based level. The rule-based level governs the choice of control alternatives. The display provides the operator with signs that he uses as cues for the selection of an appropriate action. Typically, the action alternatives consist of a set comprised of operating procedures and routine control strategies. As discussed before, the problem with conventional interfaces is that the cues they provide the operators with are not uniquely defining with respect to the current process state. The result is that the cues that operators rely on are optimized for frequently encountered situations, but they can lead to 'procedural traps' in novel situations. EID attempts to overcome this difficulty by developing a unique and consistent mapping between the symbols that govern the behavior of the process, and the signs, or cues, that the interface displays. This will reduce the frequency of errors due to procedural traps because the cues for action, being based on abstract process properties, will be uniquely defining with respect to the underlying system state.

8.1.3. Knowledge-based level. Knowledge-based behavior consists of abstract reasoning based on a mental model of the process. EID supports this level of cognitive control through the mapping of signs onto symbols. This mapping turns out to be very complex because the symbolic reference can be to several different conceptual levels representing general functions, physical processes, or equipment anatomy, depending on the actual circumstances (Rasmussen and Goodstein, in press). This means that, in addition to serving as cues for action, the same display configuration can also be interpreted in several ways as symbols for reasoning. Thus, if the display configuration is interpreted symbolically, it presents the operator with a visible model of the process that can support thought experiments and other planning activities. In addition, it is suggested that such a mapping will also support functional understanding necessary for error recovery. If signs can also be interpreted as symbols, then this may force the user to consider informative aspects when looking for action cues (Rasmussen and Vicente, 1987).

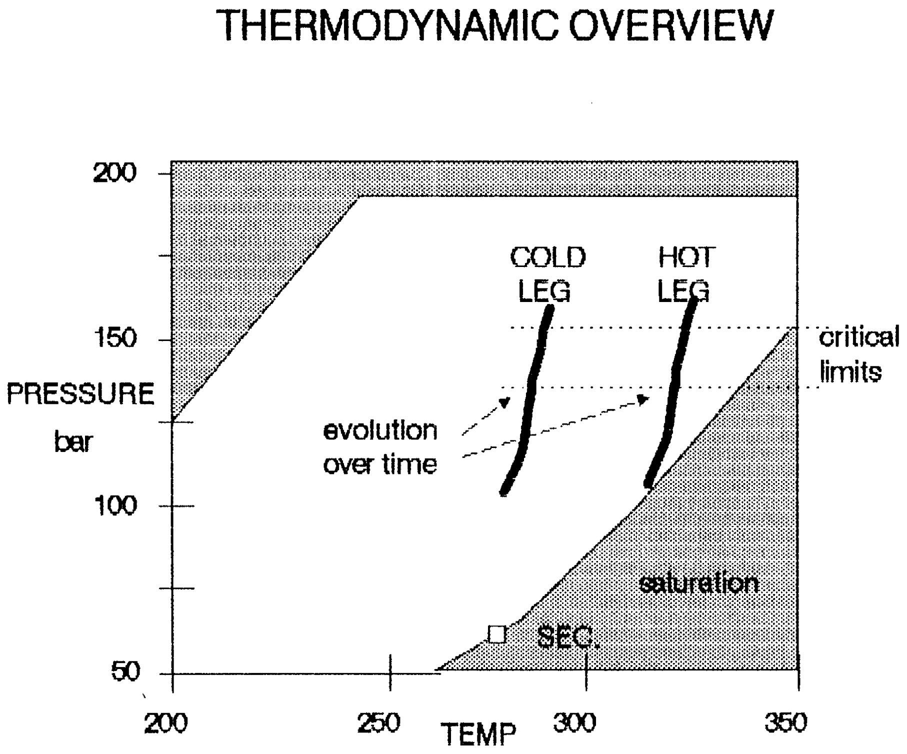

Figure 13 is a display intended to support the operators' disturbance handling and speculations about energy and mass conditions by including a considerable amount of information which has to do with the generic functioning of a power plant i.e. in this case, the thermodynamics of the process. It has the form of a P(ressure)-T(emperature) diagram indicating primary and secondary P-T trajectories vs. time.

Display formats having these features can, in some cases, readily be developed from the 'externalised mental models' which are normally being used as a support of functional reasoning in the form of graphic representation of relational structures such as technical drawings, graphs, and diagrams. Semiotic analyses of the use of such professional representations in actual work (Cuny and Boyé, 1981) have shown that they are actually interpreted as prescriptive signs or descriptive symbols, depending on the requirements of the task. Such an interface based on the engineering representation of two-phase thermodynamic systems in terms of a Rankine cycle diagram has been proposed by Beltracchi (1987).

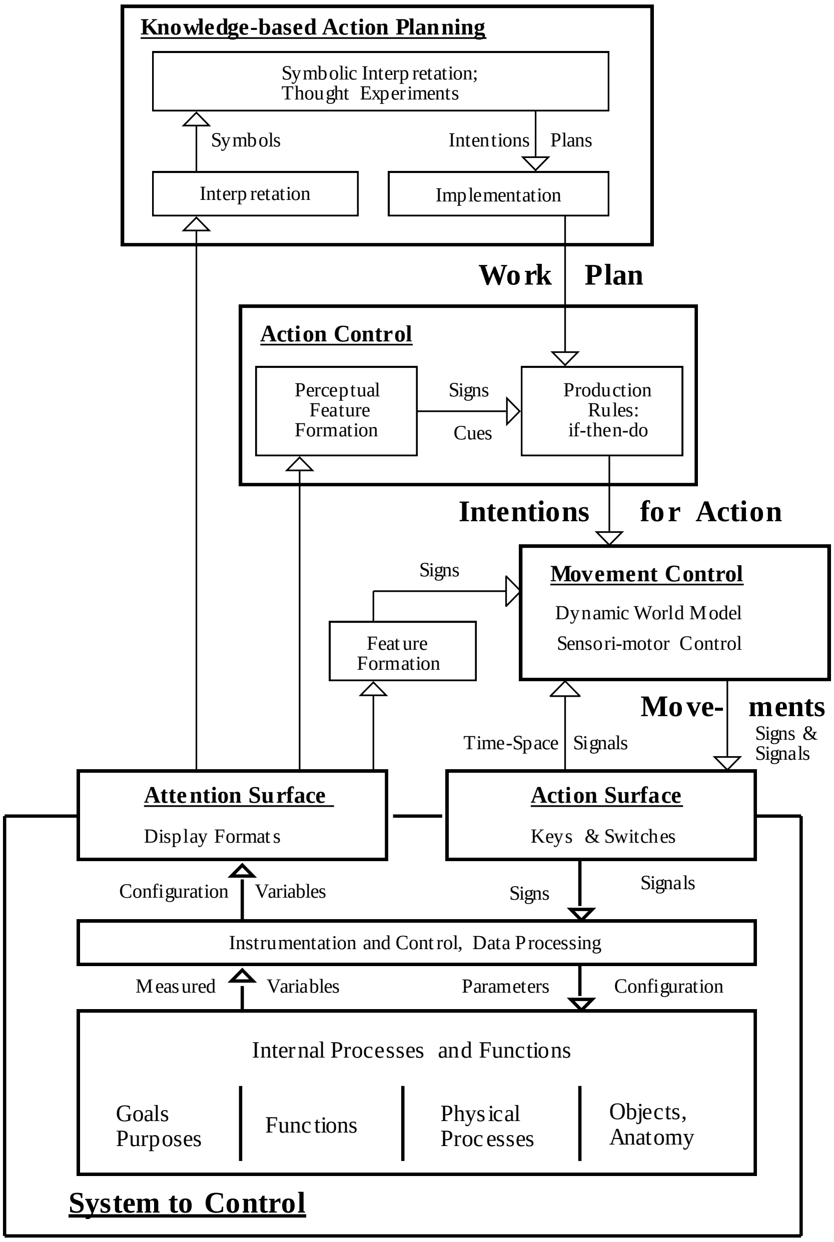

Figure 14. The figure illustrates the complex mapping between the different levels of representation of the invisible process and the different levels of cognitive control of operator action.

8.2. Musical Skill

It will be useful to analyse the perceptual and motor mappings in activities where high capacity manual skills are involved. Musical performance is such a direct manipulation activity, requiring a very high information processing capacity in a real-time, on-line mode of control. Another particular feature of this task is that the manipulation surface and the information surface are separate, as is typically the case in conventional human-machine interfaces.

Musical performance is based on a repertoire of highly automated and coordinated patterns of movements. The musical notation in a score is perceived in terms of signs referring to such complex patterns. In a text book for playing an instrument (e.g., a recorder), a table of fingering patterns are given for the individual notes. From such instructions, one can learn to play an instrument without any knowledge about the symbolic meaning of the notation. When high performance skill evolves with practice, increasingly complex patterns of notes are related to integrated patterns of movements resulting in a very high capacity and speed in performance. This re-coding requires a notation in which chunking of the individual signs occurs in a way which maps directly onto a concurrent chunking of the required movements.

The structure active for the aggregation in the sign domain should map onto the structure useful for chunking patterns of movements. Chunking of movements involves operation in a part-whole relationship and, consequently, this structure should be reflected in the notation. Organisation of movements should be reflected directly in the spatial, graphic organisation of the musical representation. The important feature of the present musical notation system is that the perception of patterns of signs can be changed concurrently with the evolution of higher level motor patterns. The mapping between perceptual and motor patterns is not a simple one-to-one mapping, but can be changed dynamically by higher level signs indicating key-changes, rhythmic instructions, musical style indications, etc. Thus, the internal dynamic world model (i.e., the attunement of the organism) can be modulated at will by the performer. The direct analogical mapping from visual patterns to motor patterns is such that changes in the part-whole dimension in both domains can easily and dynamically be performed according to level of practice, and to higher level performance indicators, e.g., instructions from the conductor.

It is possible to design more informationally 'economic' notation systems which are not directly structured in a part-whole hierarchy, but according to a generic tree in combinatorial coding instead. In that case, the transformation to action patterns would require analytical, rather than perceptual, re-coding, which in turn would make it more difficult to adapt the coding according to practice and style. Such a system has, in fact, been proposed by Jean-Jacques Rousseau (1742) in a presentation for the Academy of Sciences in Paris. His system was based on a logical notation in terms of a number code. It was well accepted by the scientists of the Academy, and a committee was founded to review the system for recommendation. The arguments of this committee mostly considered whether the system was new and therefore useful. Apparently, numerical systems for musical notation were subject to discussion at that time. However, as soon as Rousseau presented his system to the composer Rameau, he was met with arguments that convinced him: "Your signs are excellent, with respect to representation of tone and interval ---; but they are very poor because they require an activity of thought which cannot keep pace with the performance. The location of the signs in our usual notation imprint on the eye without support of this kind. If two notes, one very high and another very low are connected with a sequence intermediate notes, I immediately by first glance perceive the gradual rise from one to another. By your system, however, I necessarily have to spell my way through from number to number, one glance will not do it." Rameau's arguments fit very well the discussion presented above.

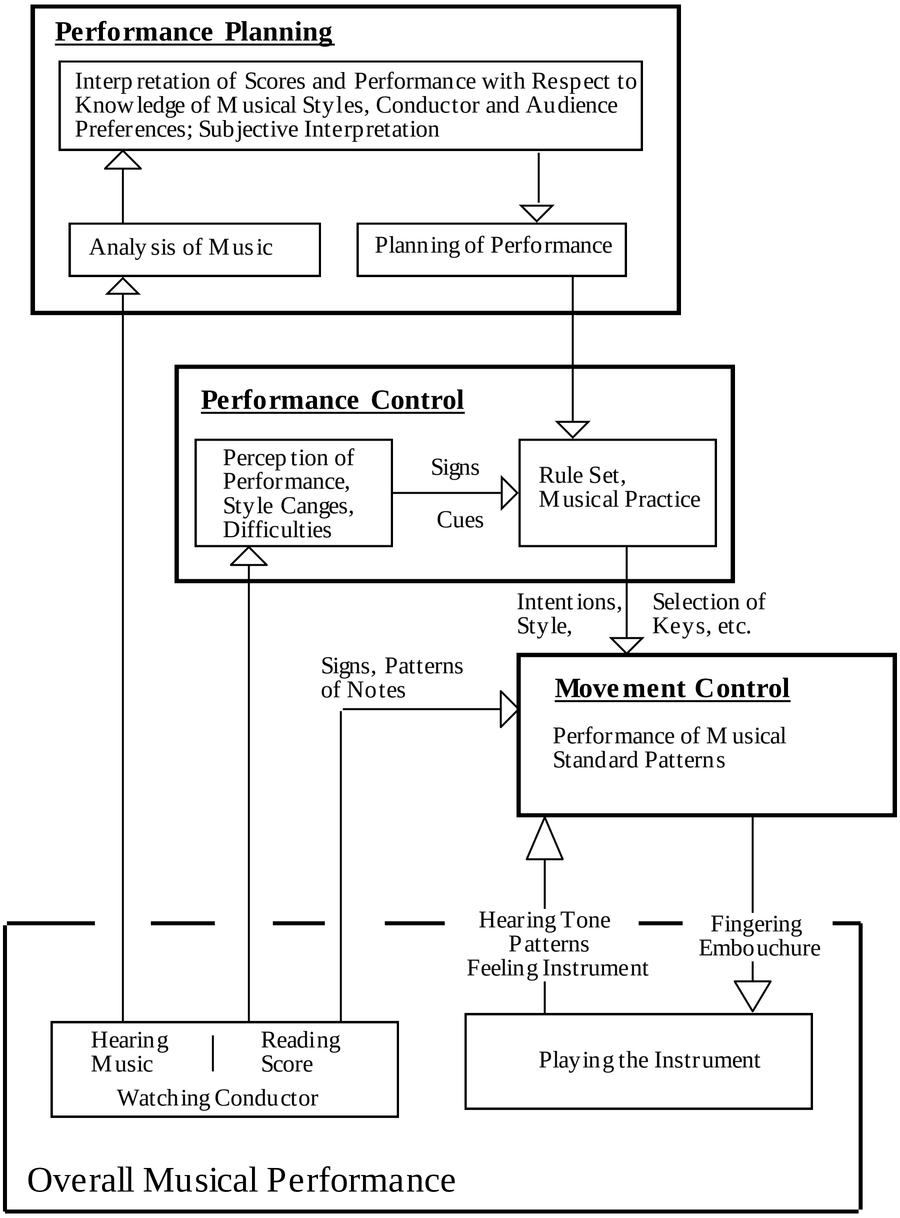

Figure 15. The figure illustrates the interaction between control of movements, control of actions, and planning of actions during musical performance.

A point worthy of further attention is the 'direct manipulation' nature of musical performance, in spite of the fact that the visual presentation of the score and the manipulation surface are separated. The determining issue is probably related to the fact that the music attended to and the score perceived visually have the same part-whole structure and merge into one perceived pattern which maps directly onto the patterns of movements - one 'sees' the music in the score. This, in fact, implies that, for this particular domain, the attention and manipulation 'surfaces' merge at the level of the music.

Another example similar to musical performance is the high performance skill developed by users of the Abacus of China and Japan. It is well known that for the elementary arithmetic operations, the speed and accuracy of skilled Abacusians are comparable to that of calculator users. The trick appears to be the same as in musical performance, namely that the numbers to introduce in a calculation are read in terms of actions on the abacus, while the arithmetic operations signs + or - act like the transposition signs in music to select the number-action mapping set needed in the particular operation. It is interesting to see that instruction books for abacus users stress the need to be consistent in the use of the proposed number-action mapping; even simpler actions are possible in some cases, in order to develop a high capacity sensorimotor skill.

8.3. Bibliographic Search

Bibliographic search is characterized by a lack of coherent internal structure of the type possessed by physical systems. The items of the database are not interrelated by causal laws like the components of a process plant, nor is there a direct mapping between the action patterns required for interaction with the computer and the feedback from the system at the semantic information level, as was the case in the musical example. In order to obtain such a mapping, special precautions should be taken to create an internal structure in the database which is homomorphic with the user's problem space. The system should possess an internal structure that will present a space in which the user can navigate easily by means of the search questions which naturally emerge from the current problem, in contrast to navigation guided by knowledge about the computer system. Typically, the internal structure of a data base is chosen to make location and retrieval of information items unambiguous, effective, and fast, i.e., it is based on a formal generic hierarchy (cf. Smith and Smith, 1977). For effective access by a user, however, retrieval should be possible from several different points of view. In a means-end space for instance, items should be accessible by question of 'what' they are, as well as 'why' and 'how' they should be used.

Sensorimotor control of movement patterns is concerned only with manipulation of items on the visible surface of a system. The advantage of the mouse interface for this task is that the communication of spatial-temporal aspects in the perception-action loop is intact. The commands sent to the computer are selected from a repertoire presented on the screen. They are identified by means of their physical position, and pointed at with the mouse before the selection order is transmitted by a 'click' sign. This implies a direct relationship between the movement pattern and the perceptual control, as was the case in musical performance. Consequently, very efficient navigation in a database would be possible if the abstract attributes of the database items could be consistently mapped onto positions in a spatial representation. In this way, location in the database could be identified perceptually, while navigation in the space could be analogically controlled by patterns of movements. This type of design has been systematically used in spatially structured databases on video disks, in which the exploration of the database is done in term of 'wandering around' via a joystick.

For support of knowledge-based reasoning, however, there are several problems. First, the number of abstract search dimensions, as defined by the dimensions of the relational structure of the user's mental model, is quite high. Secondly, the number of items in the database is large. Finally, the attributes for search are not immediately and explicitly known by the users who, therefore, may want to 'browse'. One way to exploit the capabilities of computers for flexible information presentation would be to relate the information items to the location in a virtual space, a store house, in which three dimensions of the multidimensional attribute space are represented by the location in a room, while the remaining relevant dimensions are taken care of by arranging for several rooms and departments in the store. Computers systems have the advantage that the same item can be found in several locations according to different search attributes. Thus, the stock can be browsed according to location when the rooms representing the most relevant dimensions are visited. In this way, it will be possible to take advantage of the mnemonic power of the Method of Loci (cf. Fuller, 1898) by applying it to a multi-dimensional representation. In other words, to use the idea of George Miller (1968): Information is a question of 'where'. Visual presentation of the space and analogical control of movements (mouse or joystick) will result in skill-based control of the search itself. In this way, a direct relation between movements and the location in the information space is preserved.

The structure of commands and the organisation of information presentation could be suggested from analogy to musical performance. A basic hypothesis to suggest, based on the musical analogy, is that novices are focusing on the manipulation surface, the 'tool handle', while skilled performers are focusing on the performance content. To enable an easy shift in the degree of chunking of action elements and perceptual elements as a function of the general level of training and the current task demands, the structure of the perceptual and the motor representations should map one-to-one along the part-whole dimension. In this way, the search is formulated as an exploration of a multi-dimensional space, the representation of which depends on the strategy selected by the user.

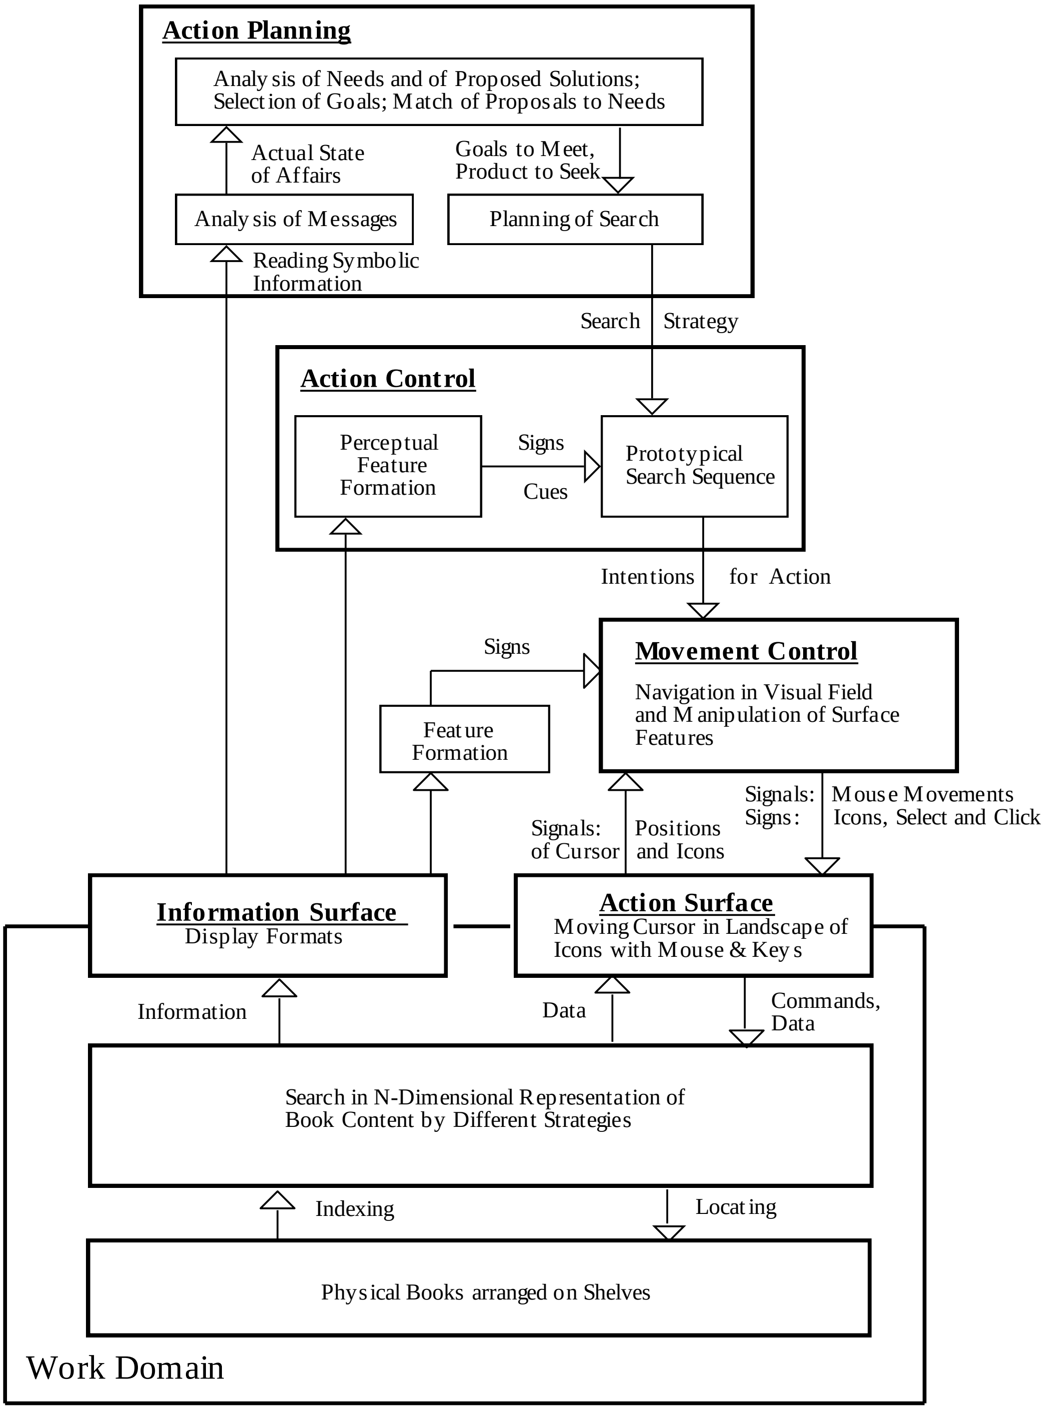

Figure 16. The figure illustrates the interaction between control of movements, control of actions, and planning of actions during bibliographic search.

An important aspect to consider is that the search will be guided by the intuition offered by the user's 'world model', even for knowledge-based control of search. One may want to consult familiar regions, to browse, or a more rational, analytical approach may be chosen. An analytical approach is a kind of selective 'addressed' search in a domain selected from a display of a 'work-rooms' showing the landscape of topical items which can be selected. In the analogy of a city map, one has an indication of the street one wants to visit, together with a helicopter which will bring one there directly. Being there, one will look at the houses. In the 'browsing strategy', the user doesn't know the address, but will recognise what he is looking for when he sees it. This means that he has to pass through the streets of the city until he recognises an item. In a library, this implies scanning the physical books on the shelf. In a computerised system, however, the user can choose a subspace of, or a 'channel' through the multi-dimensional space, by specifying some aspects of the target which are known (those that are to be selected or avoided). Thus, the difference between analytical search and browsing in a computerised system is a question of degree, not of categories.

In conclusion, for design of bibliographic search systems, it is necessary that the database reflect the user's needs in terms of search attributes. If this is not the case, a user will simply not be able to match the characteristics of unfamiliar items with his own personal needs. This implies that the problem is not just one of designing interfaces to databases. In addition, the primary information indexing must be based on intimate knowledge of the users' point of view of the actual situation.

9. AN OVERVIEW OF EID

Let us summarize the development of our ideas up to this point. We began by describing two accounts of DMI and by pointing out their limitations with respect to complex systems. In an attempt to develop a theoretical framework for interface design that would overcome some of these limitations, we described a model of HIP and extracted some implications for interface design. We then extended and refined these implications into a specific goal for interface design, couched within the concepts of the SRK framework. A morphological analysis of the transformations necessary to match the interface representation to the different levels of cognitive control enabled us to develop a set of prescriptive design principles that constitute the EID framework. In this section, we will provide a general and informal description of what EID attempts to do, as well as pointing out some of its limitations.

9.1. What It Attempts to Do

An interface is a media for communication between operator and computer. In order for communication between agents to be effective, the agents must possess a shared representation of their domain of discourse (cf. Brehmer, 1986; Cussins, 1987; Roth, Woods, Elm, and Gallagher, 1987; Vicente, 1988). Within this perspective, the goal of EID is to produce a shared perceptual representation of the problem solving world that will facilitate effective communication between man and machine. This goal can only be achieved if the representation is both normative and psychologically valid, i.e. , it must describe proper system functioning in a way that is compatible with the properties of human cognition. With its ecological roots, EID recognizes the importance of this duality between environment (task) and organism (operator).

While the communication aspect is without a doubt important, there is another primary facet of EID which we have yet to discuss. Fundamental to the EID approach to interface design is the notion that the problem solving world is bounded and described by a set of invariants. In other words, the task domain is a closed world consisting of relationships, or constraints, which always hold. The problem solving world is bounded because the purposes of the system are fixed, as are the means available for achieving those purposes. The system can be described by invariants because its behavior follows certain physical laws (e.g., conservation of mass and energy).

EID attempts to take advantage of these two factors by exploiting the task constraints so that the human operator needs less knowledge to effectively control the system. More specifically, by mapping symbols onto signs, EID is embedding the semantics of the task in the perceptual characteristics of the interface. This means that, in the limit, the operator could perform the task by acting as a symbol manipulator. Instead of relying on symbolic reasoning, he could control the system via the perceptual characteristics of the interface alone. Embodying the domain semantics in the interface would allow an operator that did not possess the theoretical knowledge of how the process functioned to exhibit the skilled level of performance that would only be expected of someone controlling the system from the fundamental principles that govern its behavior (cf. Cuny and Boye, 1981; Roth et al., 1987). Or to put it more succinctly, EID allows operators to simulate deep control by relying on surface control.

An example will help to make these points clear. The following passage was originally intended to describe the Situated approach to language being developed at Stanford (cf. Barwise and Perry, 1983). However, because of the conceptual parallels between situation semantics and EID, it will serve our purposes equally well.

> "According to the Situated theory of the user interface for video cassette recorders, the user need only represent to himself the sequence of button presses in order to successfully operate the VCR. He need know nothing about what effects those button presses are causing in the innards of the machine, because the task domain for Sony VCR operations is reasonably closed and tightly structured .... The point is

> that the Situated VCR user is a formal symbol-manipulator and not an intelligent agent who operates the VCR on the basis of his understanding of how VCRs work. Notice also that this

> symbol manipulation ability does not generalize .... What works for a Sony doesn't work for a Philips" (Cussins, 1987, p. 16).

Several caveats are in order. First, some readers will no doubt object that providing the operator with the role of a mere ‚symbol-manipulator' is not very satisfying. To this argument, we have two replies. First, the idea of the operator meeting the entire range of task demands by acting as a symbol-manipulator (i.e., by relying merely on RBB and SBB) is an unattainable ideal. There will always be circumstances where the operator will resort to KBB. Our point, however, is that interfaces should be designed to support symbol-manipulating behavior so that the operator can effectively rely on it when he wishes. As we mentioned previously, this is an adaptive strategy since lower levels of cognitive control are generally less effortful than higher levels. Secondly, while the symbol-manipulating role may sound demeaning, the fact of the matter is that that is precisely what operators are already doing with existing systems. As discussed earlier, operators have shown a consistent and overwhelming preference for lower levels of cognitive control.

A second possible objection to the EID approach is: What happens when there is a failure? It could be argued that everything is fine under normal circumstances but when something unexpected happens the same procedures no longer work. The VCR example is somewhat misleading in this respect since there is not a 1:1 mapping between signs and symbols, i.e., the surface features do not map uniquely onto the fundamental semantic properties. Thus, if there were a breakdown, the Situated VCR operator would be helpless. However, in the case of an EID interface which is based on invariants, the operator can also cope with faults. The main reason behind this is that the invariants always hold. Even when a valve is blocked, the laws of conservation of mass and energy still describe the behavior of the system. This means that faults can be detected as deviations from proper system functioning (Rasmussen and Lind, 1981), thereby allowing the operator to recover effectively from unexpected incidents.

9.2. What It Doesn't Do

In order to fully understand EID, it is important that we discuss the scope of the approach as well as some of its apparent limitations. In terms of applicability, the EID approach to interface design can only be applied to problem solving domains that are closed and that have well-defined functions and processes. The reason for this is that an EID interface has built-in transformations that make these functions and processes visible to the user. In cases where the functions relevant to the user do not constitute a stable, well defined set, it will not be possible for the designer to specify a normative representation that will be appropriate for the entire range of task demands that the operator will encounter. In most cases, this limitation does not present a problem since high technology installations are designed with very specific purposes in mind.

Secondly, because EID attempts to take advantage of the constraints specific to the problem domain, the competencies that operators acquire with a certain system will not transfer to other kinds of systems. Again, this is not very problematic since, in the majority of cases where information technology is applied to the work domain, each agent operates wihin the same work context day after day. For instance, in the case of process control, the operator always controls the same process. In addition, the advantage of information technology is its flexibility that allow interfaces to be tailored to the individual work domains.

It follows that EID requires that the designer have a good understanding of the problem domain. In particular, the designer must know what the invariant relationships that describe the behavior of the system are. In designing interfaces for systems which require diagnosis and disturbance control, the designer must carefully analyse the various categories of faults in a manner similar to that usually followed for risk and safety analysis. As mentioned above, EID tries to take advantage of the constraints inherent in the domain. Thus, it is imperative that the designer know exactly what these constraints are if he is to embed them in the interface. Due to rapid advances in the modelling of physical systems, there are many high technology systems which EID can be applied to. Nevertheless, it is important to point out that there are facilities being constructed whose behavior is not well understood by their designers. In these situations, the application of EID is problematic at best.

The limitations discussed to this point have to do with the applicability of the EID framework. We hope to have shown that the EID framework has the potential to be successfully applied to a wide variety of systems. The only factor limiting the generalizability of the approach is the designer's knowledge or understanding of the system being controlled.

To summarize, EID attempts to reveal the complexity of a work domain to the operator in a f orm that he can cope with. As Newman (cited in Rasmussen, 1974, p. 27) states: "People don't mind dealing with complexity if they have some way of controlling it--- (If) a person is allowed to structure a complex situation according to his perceptual and conceptual needs, sheer complexity is no bar to effective performance". This is critical to the understanding of EID. The approach does not relieve the operator of the task demands as an automated control algorithm would. The line of attack is a different one. Given a fixed set of task demands, EID attempts to answer the question: What is the best way of presenting the complexity inherent in the domain? This question looks at the problem from the side of the task. The dual question, arising from the side of the operator is: What is the most effective resource the operator has to deal with complexity?

10. WHAT'S IN A NAME?

The name ecological interface design alludes to a philosophy that is a very important part of our framework. We are referring to the ecological approach to psychology first advocated by Brunswik, and later, albeit in a different form, by Gibson (cf. Brunswik, 1957 and Gibson, 1966; 1979). In this section, we will try to relate EID to each of these approaches.

10.1. Probabilistic Functionalism

The basic premise of Brunswik's theory of probabilistic functionalism is that psychology should be concerned, not just with the organism, but more importantly, with the interrelationships between the organism and its environment (Brunswik, 1957). Thus, Brunswik distinguished between two types of stimulation: distal variables represent objective descriptions of the state of the organism's ecology, whereas proximal variables represent the sensory input that the organism receives from its ecology. Brunswik believed that the organism is not able to perceive the distal variables directly, but instead must infer what is going on in the ecology from the imperfect (i.e., probabilistic) cues provided by the proximal variables. This leads to another important distinction, that between a cue's validity and its utilization. Cue validity is given by the correlation between the proximal and the distal variables. However, utilization of the sensory input, may or may not be appropriate. Therefore, the concept of cue utilization is needed to describe how the organism makes use of the cues available to her.

Given this framework, it follows that an appropriate goal for interface design would be to provide the operator with cues that have perfect ecological validity. In fact, this is the goal behind EID: by mapping symbols onto signals, we are in fact mapping distal variables onto proximal variables. Ideally, this would lead to a completely transparent system; the interface should completely and unambiguously define the current system state. Note that this does not necessarily guarantee that the organism's utilization of the cues will be optimal.

10.2. Direct Perception

Gibson's approach differed from Brunswik's in that he believed that perception was direct, i.e., that people directly perceived the higher order variables that the ecology had to offer them, without any mediating information processing. These higher order variables are combinations of the simple variables that Brunswik dealt with, and according to the theory of direct perception, they completely specify the distal objects, thereby eliminating the need for inference and probabilism (Brehmer, 1984).

Gibson (1979) introduced a new vocabulary to explain his theory of direct perception. The basis for perception is said to be the invariant relationships in the ecology that are made available to the observer via invariants in the optical array. The notion of an affordance, an invariant combination of variables that demands or invites appropriate behaviors, was introduced to account for goal-oriented behavior. Basically, an affordance represents attributes of the environment that are relevant to the organism's purposes; it specifies a possibility for action. An object's affordances are perceived via the invariants in the optical array through a process of direct attunement which is closely related to the conditioning of the neural system as represented by the internal dynamic world model underlying skill-based performance (Rasmussen, 1986, p. 79). Thus, perception is viewed as a means of selecting the appropriate action to attain a goal, and the concept of an affordance relates perception to action. The result is goal-oriented behavior.

In our case, the ecology is the process being controlled, and its invariants are described by a set of mathematical equations. Because the process is invisible, the information intrinsic in these invariants is normally not available to the operator. EID attempts to map the invariants in the process onto invariants in the interface. Again, the idea is to make visible the invisible. According to Gibsonian theory, these invariants in the interface should allow operators to directly perceive the system's affordances. Thus, EID can be viewed as building into the interface the affordances that the operator needs to effectively control the system. Because the system is best described in terms of an abstraction hierarchy (Rasmussen, 1986), the process will actually be described in terms of a hierarchy of higher-order invariants at various levels of abstraction.

This description of EID within the framework of direct perception implies that the system could be controlled without any mediating decision making; the information in the invariants would be perceived as affordances that would specify what action to take. This is only true if we are rigorous in our extension of Gibson's concept of affordance to the domain of complex systems. While in certain cases, it is possible to design an interface that will offer an affordance in the classical sense of a possibility for action, we will argue that there will be other cases where this will not be possible.

Using Rasmussen's (1974; 1986) decision ladder as a framework, we see that there are three general states of knowledge that an interface can provide an operator with. These are:

1. Current system state;
2. Target state to be achieved;

3. Action to carry out.

Each successive category requires more information in order to remove the degrees of freedom that the operator is left with. The third category is the case of direct perception. But one of the fundamental assumptions of the theory of direct perception is that there is usually enough information available in the optical array to make the basic affordances of the terrestrial environment directly perceivable (Gibson, 1977). If this assumption does not hold, direct, veridical perception is impossible. We believe that, in fact, this assumption does not always hold in process systems. The natural environment is more rich in information than an industrial process.

The complexitity of process sytems means that sometimes, it is only possible to provide information at the first or second level. In these cases, the operator is required to reason in order to deal with the degrees of freedom he has available to him. For instance, because of the many-to-many mapping between levels in the abstraction hierarchy, there are various ways to attain a given target state. This creates a need for planning since the appropriate path to take cannot be determined from the available information. In addition, the fact that each fault is unique means that, even if the current system state is completely specified, the target state to be achieved is not obvious. Again there are degrees of freedom, this time in determining, for instance, whether it is better to keep the system on-line and try to compensate for the fault, or whether to suspend production. Thus, there is a need to evaluate conflicting criteria before deciding what the goal state should be. It should also be noted that, due to uncertainty in the data obtained from sensors, it is even sometimes difficult to accurately determine what the current system state is. To summarize, depending upon the circumstances, control of the system will be more or less mediated by inferencing.

Reinterpreting Gibson's theory in terms of the SRK framework will show how EID is related to direct perception. In the natural environment, there is enough intrinsic information in the optical array for interaction to take place at the skill-based level. This represents the case of affordances in terms of actions (category 3 above). However, the complex nature of process systems implies that the operator will have to resort to higher levels of cognitive control at certain times. This will occur when he is only afforded the target state to achieve (category 2 above), or the current system state (category 1 above). During these times, rule- or even knowledge-based behavior will be necessary to successfully cope with task demands. In order to deal with these situations, EID attempts to make the most of the information that is available in order to afford the operator as much as possible. It is only because of the comparative complexity of process systems that control via direct perception is not possible as it is in the natural ecology.

11. DIRECT MANIPULATION REVISITED

As mentioned at the outset, the goal of this paper was to develop a theoretical framework that overcame the limitations of current accounts of DMI. We are now in a position to summarize how far EID has come beyond DMI.

First, our theoretical framework has allowed us to decribe the benefits of DMI in terms of properties of the HIP system. This description is both more rigorous and more general than Shneiderman's SSM and the theory provided by Hutchins et al. Secondly, as the name suggests, most discussions of DMI have been centered around the skill- and rule-based levels. The focus has been on displaying the domain objects of interest and allowing the user to act directly on those objects. However, displaying semantically rich objects is much more difficult to do in process control because there is a very complex mapping between the various levels in the abstraction hierarchy (Rasmussen and Vicente, 1987). In addition, the domains that DMI have been applied to (e. g., text editing and spreadsheet calculations) are simple enough that the user can infer what the underlying relational structures are without having them explicitly represented. We have argued that this is not the case in process control, where functional reasoning takes place in a complex, causal network. Therefore, in order to deal with the complexities of process control, an interface must provide support for all three levels of cognitive control. EID provides some principles on how this should be done.

12. CONCLUSION

It is fitting that we end a paper of a theoretical nature on a philosophical note. In their discussion on DMI, Hutchins et al. (1986, pp.119-120) make the following comment.

> "On the surface, the fundamental idea of Direct Manipulation interface to task flies in the face of two thousand years of development of abstract formalisms as a means of understanding and controlling the world. So, the exterior of Direct Manipulation, providing as it does for the direct control of a specific task world, seems somehow atavistic, a return to concrete thinking. On the inside, of course, the implementation of direct manipulation systems is yet another step in that long formal tradition. The illusion of the absolutely concrete world is made possible by the technology of abstraction".

We agree with Hutchins et al. that DMI are contrary to the tradition of formal thinking. Given our ecological perspective, we would even go so far as to say that that is the whole point: people excel at utilizing the effortless, parallel, routines available to subconscious processing, not at abstract, formal reasoning. EID recognizes this fact and provides guidelines for how to support people in what they do best. Knowledge-based processing should only be activated when the demands of the domain require it. Yet even then, EID recognizes that abstract reasoning is an effortful and error prone process, and thus, it attempts to provide the support necessary to carry out such activities effectively.

The fact that the illusion of a concrete world is only made possible through the application of computer technology is not what is important. What is important is that systems designed in the EID philosophy will allow people to take advantage of the perceptual-motor skills that have been perfected by millions of years of evolution, while at the same time, supporting the more recently evolved, limited-capacity reasoning activities . The end result should be the creation of interfaces that reveal the ecology of the problem solving domain in a manner that is consonant with the properties of the human organism.

ACKNOWLEDGEMENTS

This report was written while KJV was a visiting scientist at Riso. The work described here represents the theoretical framework that will serve as the basis for his doctoral dissertation research, to be conducted at the University of Illinois at Urbana-Champaign. KJV would like to thank: Jens Rasmussen for his intellectual stimulation and for being generous enough to openly share his great wealth of knowledge; and Tine, Jan, Joâo, Lisbeth, and Sandra for their deep friendship. Collectively, they turned a year in Denmark into an experience of a lifetime.

REFERENCES

Beltracchi, L. (1987). A direct manipulation interface for water-based rankine cycle heat engines. IEEE Systems, Man, and Cybernetics, SMC-17, 478-487.

Brehmer, B. (1986). Man as an operator of systems. In H. E. Peterson and W. Schneider (Eds.), Human-computer communications in health care (pp. 7-15). New York: Elsevier.

Brehmer, B. (1984). Brunswikian psychology for the 1990`s. In K. M. J. Lagerspetz and P. Niemi (Eds.), Psychology in the 1990's (pp. 383-398). New York: North-Holland.

Brunswik, E. (1957). Scope and aspects of the cognitive problem. In H. Gruber, R. Jessor, and K. R. Hammond (Eds.), Contemporary approaches to cognition (pp. 5-31). Cambridge: Harvard University Press.

Carello, C., Turvey, M. T., Kugler, P. N., and Shaw, R. E. (1984). Inadequacies of the computer metaphor. In M. S. Gazzaniga (Ed.), Handbook of cognitive neuroscience (pp. 229-248). New York: Plenum Press.

Chapman, D., and Agre, P. E. (1986). Abstract reasoning as emergent from concrete activity. In M. P. Georgeff and A. L. Lansky (Eds.), Reasoning about actions and plans (pp. 411-424).

Craik, K. J. W. (1943). The nature of exlanation. Cambridge: Macmillan.

Cuny, X., and Boyé, M. (1981). Analyse semiologique et apprentissage des outils-signes: L'Apprentissage du schema d'electricité. Communications, 33, 103-140.

Cussins, A. (1987). Being situated versus being embedded. CSLI Monthly, 2(7), 14-20.

Dreyfus, H. L., and Dreyfus, S. E. (1986). Mind over machine: The power of human intuition and expertise in the era of the computer. Oxford: Basil Blackwell.

Fuller, H. H. (1898). The art of memory. St.Paul, Minnesota: National Publishing Company.

Gibson, J. J. (1979). An ecological approach to visual perception. Boston: Houghton-Mifflin.

Gibson, J. J. (1977). The theory of affordances. In R. Shaw and J. Bransford (Eds.), Perceiving, acting, and knowing (pp. 67-83). Hillsdale, NJ: LEA.

Gibson, J. J. (1966). The senses considered as perceptual systems. Boston: Houghton-Mifflin.

Hammond, K. R., Hamm, R. M., Grassia, J., and Pearson, T. (1987). Direct comparison of intuitive and analytical cognition in expert judgement. IEEE Transactions on Systems, Man, and Cybernetics, SMC-17, 753-770.

Hollnagel, E. (1981). Report from the third NKA/KRU-experiment: The performance of control engineers in the surveillance of a complex process (NKA/KRU-P2(81)36). Roskilde, Denmark: Riso National Laboratory, Electronics Department.

Hutchins, E. L., Hollan, J. D., and Norman, D. A. (1986). Direct manipulation interfaces. In D. A. Norman and S. W. Draper (Eds.), User centered system design: New perspectives on human-computer interaction (pp. 87-124). Hillsdale, NJ: LEA.

Klein, G. A. (in press). Recognition-primed decisions. In W. B. Rouse (Ed.), Advances in man-machine systems research, vol. 5. Greenwich, CT: JAI Press.

Miller, G. A. (1968). Psychology and information. Am. Docum., 19, 286-289.

Norman, D. A. (1986). Cognitive engineering. In D. A. Norman and S. W. Draper (Eds.), User centered system design: New perspectives on human-computer interaction (pp. 31-61). Hillsdale, NJ: LEA.

Pew, R. W., and Baron, S. (1983). Perspectives on human performance modelling. In G. Johannsen and J. E. Rijnsdorp (Eds.), Analysis, design, and evaluation of man-machine systems (pp. 1-14). Oxford: Pergamon.

Rasmussen, J. (1987). Cognitive control and human error mechanisms. In J. Rasmussen, K. Duncan, and J. Leplat (Eds.), New technology and human error (pp. 53-61). New York: Wiley.

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

Rasmussen, J. (1983). Skills rules, knowledge: Signals, signs, and symbols and other distinctions in human performance models. IEEE Transactions on Systems, Man, and Cybernetics, SMC-13, 257-267.

Rasmussen, J. (1974). The human data processor as a system component: Bits and pieces of a model (Riso-M-1722). Roskilde, Denmark: Danish Atomic Energy Commission.

Rasmussen, J., and Goodstein, L. P. (in press). Information technology and work. In M. Helander (Ed.), Handbook of human-computer interaction. New York: North Holland.

Rasmussen, J., and Vicente, K. J. (1987). Cognitive control of human activities and errors: Implications for ecological interface design (Riso-M-2660). Roskilde, Denmark: Riso National Laboratory, Department of Computer and Information Science.

Reason, J. (1988). Framework models of human performance and error: A consumer guide. In L. P. Goodstein, H. B. Andersen, and S. E. Olsen (Eds.), Tasks, errors, and mental models: A festschrift to celebrate the 60th birthday of Professor Jens Rasmussen (pp. 35-49). London: Taylor and Francis.

Reason, J. (1987). Modelling the basic error tendencies of human operators. Paper presented at Post-SMIRT 9 Seminar on Accident Sequence Modelling: Human Actions, System Response, and Intelligent Decision Support. Munich, August 24-25.

Reason, J. (1984). Lapses of attention in everyday life. In R. Parasuraman and D. R. Davies (Eds.), Varieties of attention (pp. 515-549). New York: Academic Press.

Roth, E. M., Woods, D. D., Elm, W. C., and Gallagher, J. M. (1987). Providing on-line advice for a dynamic control task: A case study in intelligent decisio support system design. Proceedings of the Human Factors Society 31st Annual Meeting (pp. 36-40). Sanata Monica, CA: HFS.

Rouse, W. B. (1983). Models of human problem solving: Detection, diagnosis, and compensation for system failures. Automatica, 19, 613-625.

Rousseau, J. J. (1742). Confessions, Vol.2, Book 7.

Runeson, S. (1977). On the possibility of "smart" perceptual mechanisms. Scandinavian Journal of Psychology, 18, 172-179.

Sanderson, P.M., and Harwood, K. (1988). The skills, rules and knowledge classification: A discussion of its emergence and nature. In L. P. Goodstein, H. B. Andersen, and S. E. Olsen (Eds.), Tasks, errors, and mental models: A festschrift to celebrate the 60th birthday of Professor Jens Rasmussen (pp. 21-34). London: Taylor and Francis.

Searle, J. R. (1981). Minds, brains, and programs. In J. Haugeland (Ed.), Mind design (pp. 282-306). Montgomery, VT: Bradford Books.

Shneiderman, B. (1983). Direct manipulation: A step beyond programming languages. IEEE Computer, 16(8), 57-69.

Smith, J.M., and Smith, D. C. P. (1977): Database abstractions: Aggregation and generalization. ACM Transactions on Database Systems, 2, 105-133.

Woods, D. D. (1985). Visual momentum: A concept to improve the cognitive coupling of person and computer. International Journal of Man-Machine Studies, 21, 229-244.

Zwicky, F. (1967). The morphological approach to discovery, invention, research and construction. In F. Zwicky and A. G. Wilson (Eds.), New methods in thought and procedure (pp. 273-297). Berlin: Springer-Verlag.

Risø National Laboratory **Risø – M –** 2736

Title and author(s) A THEORETICAL FRAMEWORK FOR ECOLOGICAL INTERFACE DESIGN Kim J. Vicente and Jens Rasmussen	**Date** August 1988 **Department or group** Information Technology **Groups own registration number(s)** R-9-88 **Project/contract no.**
Pages 76 **Tables** 1 **Illustrations** 16 **References** 43	**ISBN** 87-550-1459-3

Abstract (Max. 2000 char.)

A theoretical framework for designing interfaces for complex systems is described. The framework, called ecological interface design (EID), suggests a set of principles for designing interfaces in a way that supports the fundamental properties of human cognition. The basis of EID is the skills, rules, knowledge model of cognitive control. In order to support the full range of operator problem solving activities, EID suggests how to design interfaces that simultaneously support each of the three levels of cognitive control, but that do not force processing to a higher level than the demands of the task require. The EID approach extends the concept of direct manipulation interfaces by taking into account the added complications introduced by complex systems. In this paper, we describe the development of the framework, its theoretical foundations, and examples of its application to various work domains.

Descriptors

Available on request from Risø Library, Risø National Laboratory, (Risø Bibliotek, Forskningscenter Risø),
P.O. Box 49, DK-4000 Roskilde, Denmark.
Telephone 02 37 12 12, ext. 2262. Telex: 43116, Telefax: 02 36 06 09