The Ecology of Human-Machine Systems II: Mediating 'Direct Perception' in Complex Work Domains

Vicente, Kim J.; Rasmussen, Jens

Published in:
Ecological Psychology

Link to article, DOI:
[10.1207/s15326969eco0203_2](https://doi.org/10.1207/s15326969eco0203_2)

Publication date:
1990

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Vicente, K. J., & Rasmussen, J. (1990). The Ecology of Human-Machine Systems II: Mediating 'Direct Perception' in Complex Work Domains. *Ecological Psychology*, 2(3), 207-249.
https://doi.org/10.1207/s15326969eco0203_2

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

THE ECOLOGY OF HUMAN-MACHINE SYSTEMS II: MEDIATING "DIRECT PERCEPTION" IN COMPLEX WORK DOMAINS¹

Kim J. Vicente,

Department of Mechanical and Industrial Engineering
University of Illinois at Urbana-Champaign

Jens Rasmussen,

Risø National Laboratory

ABSTRACT: Recently, a new class of artifacts has appeared in our environment: complex, high technology work domains. An important characteristic of such systems is that their goal-relevant properties cannot be directly observed by the unaided eye. As a result, interface design is a ubiquitous problem in the design of these work environments. Nevertheless, the problem is one which has yet to be addressed in an adequate manner. An analogy to human perceptual mechanisms suggests that a smart instrument approach to interface design is needed to supplant the rote instrument (single-sensor-single-indicator) approach that has dominated to this point. Ecological Interface Design (EID) is a theoretical framework in the smart instrument vein that postulates a set of general, prescriptive principles for design. The goal of EID is twofold: first, to reveal the affordances of the work domain through the interface in such a way as to take advantage of the powerful capabilities of perception and action; and second, to provide the appropriate computer support for the comparatively more laborious process of problem solving. An example of the application of the EID framework is presented in the context of a thermal-hydraulic system. The various steps in the design process are illustrated, showing how the abstract principles of EID can be applied in a prescriptive manner to develop a concrete design product. An important outcome of this discussion is the novel application of Rasmussen's means-end hierarchy to structure the affordances of an ecosystem. The means-end hierarchy is a generic framework for describing goal-oriented systems with many degrees of freedom, and therefore represents a useful addition to the conceptual framework of ecological psychology.

All sorts of instruments have been devised for mediating apprehension. Some optical instruments merely enhance the information that vision is ready to pick up; others ... require some inference; still others ... demand a complex chain of inferences. Some measuring instruments are closer to perception than others. (Gibson, 1979, p. 260)

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 able to improvise rapidly ... if he is allowed to break down ... 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. (Rasmussen, 1974, p. 11)

¹In: Ecological Psychology, 1990, 2(3), 207-249.

INTRODUCTION

During the last 20 years, a new class of artifacts has appeared in our environment: complex, high technology work domains. Typical examples are nuclear power plants, emergency response centers, flexible manufacturing systems, air traffic control centers, and intensive care units. The number of people who either work in these places or are exposed to their effects (both positive and negative) is considerable enough that complex human-machine systems are a significant component of our contemporary ecology. While these systems provide new technical capabilities that were never before possible, they also create a new set of problems for their operators. Evolution has provided people with the capabilities required to survive in the natural environment, but those capabilities are not necessarily the most desirable ones for coping with the demands of high technology work domains.

The field of cognitive engineering has emerged in an effort to deal with this class of applied problems (cf. Norman, 1981; Rasmussen, 1986). In very general terms, cognitive engineers are concerned with designing safe and reliable high technology systems by giving appropriate consideration to the capabilities and limitations of the human element in the system. Thus, cognitive engineering can be considered as a human factors approach that is specifically tailored to complex human-machine systems. Overviews of the field are provided by Hollnagel and Woods (1983), Norman (1986), Rasmussen and Goodstein (1988), and Woods and Roth (1988), while collections of specific studies can be found in Goodstein, Andersen, and Olsen (1988), Hollnagel, Mancini, and Woods (1986, 1988), Rasmussen, Duncan, and Leplat (1987), Rasmussen and Rouse (1981), and Sheridan and Johannsen (1976).

The subject of this paper is interface design, a particularly important cognitive engineering issue, which, perhaps surprisingly, has yet to be treated in a satisfactory manner. Several aspects of interface design for complex work domains are directly relevant to ecological psychology, and it is the aim of this paper to point these out. First, there are significant conceptual parallels between cognitive engineering and ecological theory, particularly the similarity between the study of smart perceptual instruments and the design of "smart" interfaces. Second, the problem of interface design provides a new context for testing existing ecological theories of event perception and action. Third, and perhaps most importantly, the applied problems with which cognitive engineering is concerned also pose fundamental psychological questions. These basic research issues provide a set of challenges that can serve as a catalyst for extending the scope of ecological psychology.

The remainder of the paper is organized as follows. First, a brief description of the characteristics of complex work domains will serve to outline the general context of concern to cognitive engineers. In the second section, a set of implications for interface design will be derived from an analogy to perceptual instruments. To anticipate, the conclusion will be that a smart instrument approach to interface design is required. Third, a framework for structuring the affordances of a work domain, called a means-end hierarchy, will be described. In addition, several other potential applications of the means-end hierarchy that are pertinent to ecological psychology will be suggested. The fourth section will show how the preceding arguments lead to a novel theoretical framework, called Ecological Interface Design (EID). Fifth and finally, a detailed application of the principles of EID will be presented in the context of a thermal-hydraulic process control system. The intent is to illustrate how the general principles of EID can be instantiated to develop a concrete design product.

THE PROBLEMS OF COGNITIVE ENGINEERING

Complex, High Technology Work Domains

Complex work environments, or domains, possess several characteristics which have important implications for the demands that they place on their operators (cf. Perrow, 1984; Rasmussen and Vicente, in press; Sheridan, 1987; Woods, 1988). First, these systems are dynamic and tend to have long time constants; events evolve much more slowly than in the natural environment. Because feedback is delayed, the perception-action loop cannot be continuously maintained, and operators will have to rely on a qualitatively different type of control strategy (e.g., Crossman and Cooke, 1974; Veldhuyzen and Stassen, 1976). Second, there is also a high degree of risk in complex work domains, since inappropriate control actions can have catastrophic consequences. This suggests that operators may have to evaluate carefully the consequences of their actions based on a conceptual understanding of system functioning before actually implementing those actions. This will be particularly critical in abnormal (i.e., fault) situations. Third, complex work domains also tend to be composed of many subsystems that are highly coupled. This makes it very difficult to predict all of the effects of a control action, or to trace all of the implications of a fault since there are many, perhaps diverging, propagation paths. Fourth, complex human-machine systems are highly automated. Most of the time, automated control loops regulate the system, and the operator's responsibility is to monitor the state of the system. Thus, the operator's demands are primarily perceptual and cognitive in nature, rather than

psychomotor. Fifth, there tends to be uncertainty in the data that is available to operators. Because of this impoverished input, the data from the interface does not always uniquely specify the state of the system. Quite often, there will be a need for complex problem solving (i.e., to "go beyond the information given"). Finally, if this were not enough, operators are also responsible for dealing with fault situations. It is their responsibility to improvise and adapt to the contingencies of an abnormal event, quickly to ensure system safety. Since their normal control procedures no longer apply in these cases, operators must generate an appropriate response based on a conceptual understanding of the system.

The inventory of properties listed above, and their associated implications, pose an enormous burden on operators and designers. The job of the operators is particularly unenviable, since they must cope with all these demands on-line, in real time. While there is certainly a prominent place for direct perception in the control of these systems, it is important to realize that, due to the large degree of complexity, operators will also have to rely on problem solving (Rasmussen and Vicente, 1989, in press), particularly when faults occur. In order to effectively carry out these activities, the operator must have comprehensive information about system state and accurate knowledge of system structure and functioning.

A design challenge arises because contemporary human-machine systems are qualitatively different in nature from traditional work domains. Classic work domains (those which traditional human factors practices have been tailored to) are less complex in that they do not share all of the properties listed above (Rasmussen, 1988a; Rasmussen and Vicente, in press). Because complex work domains pose a new set of challenges, new methods and principles are required for system design (Hollnagel, 1983; Rasmussen, 1988a; Woods, O'Brien, and Hanes, 1987).

Interface Design

There is another characteristic of complex systems that is of prime importance. Because of the very nature of the systems being controlled, the goal-relevant properties of the work domain usually cannot be directly observed by the unaided eye¹ (e.g., the thermodynamic heat engine cycle in a power plant). For the most part, it will not be possible for operators to go out and directly explore the status of the system using the powerful perceptual systems that serve them

¹This paper will only be concerned with visual interfaces. While there may be much to be gained by exploiting other perceptual modalities in interface design, comparatively little work has been done in this area (but see Blattner, Sumikawa, and Greenberg, 1989 and Gaver, 1986, 1989 for some interesting ideas on auditory display design). Thus, the design of multi-modal interfaces remains a topic for future research.

so well in the natural environment. As a result, the control room interface (which includes controls, displays, and alarms) must serve as operators' "window" to the work domain, providing them with a mediating representation of the system. Thus, interface design is a ubiquitous problem in the design of complex systems. Yet, it still has not been addressed in a satisfactory manner (DeBor and Swezey, 1989). There is an obvious and immediate need for a prescriptive set of principles that will guide designers in developing effective interfaces for complex work domains.

What role should an interface play in high technology systems? Since the interface serves as the basis for perception, it should provide as rich a source of information regarding system status, structure, and function as possible. Can concepts from ecological psychology help in attaining this objective?

TWO TYPES OF INSTRUMENTS

On the surface, cognitive engineering and ecological psychology may seem to be concerned with very different types of questions. What does the design of interfaces, for instance, have to do with the study of event perception and action? If one considers perception as a measurement process, then these two areas no longer seem so remote. Perhaps, understanding how perception of the natural environment takes place can lead to insights into how to design effective interfaces for complex work domains.

Rote Instruments

Given the view of perception as a measurement process, the organism's perceptual systems can be considered as a set of measurement instruments (Rosen, 1978; Runeson, 1977; Pattee, 1986; see Greene, 1982 for a related discussion on action instruments for motor control). In an influential paper, Runeson (1977) makes the distinction between rote and smart instruments. As it turns out, this distinction is as relevant to interface design as it is to perception.

Rote instruments consist of a large number of a few basic types of meters, each of which performs a rather simple task. An example is a ruler, which measures a fundamental dimension, length. The advantage of rote instruments is that they can be used to derive a variety of different properties. For example, a ruler can be used to measure length, area, and volume. The disadvantage of rote instruments is that, in order to derive more complex properties, the observer must know the rules for combining the elemental measurements. For instance, to derive the area of a triangle with a ruler, the measuring agent must know the appropriate formula. In addition, calculations are also required for

the derivation of the higher order property of interest. In the case of biological systems, these computations may require considerable effort, and in fact, the computations may exceed the organism's capabilities or the time available for effective action. Runeson (1977) argues that the metaphor of rote instruments is not a very appropriate one for perception. For a goal-directed organism, the most important properties of the environment likely will be complex, higher-order dimensions that are directly relevant to the immediate goals, not the fundamental dimensions of physics. For example, it is more important for a person to know whether a certain object affords sitting than to know the exact dimensions of the object. The suggestion is that measurement with rote instruments would be very inefficient, if not intractable. Returning to the sitting example, it would be very difficult for an observer armed only with a ruler to determine whether a given object was indeed sit-onable or not. In addition to the extensive knowledge of geometry, a very large number of calculations would also be necessary, requiring a great deal of effort and time. A rote instrument approach to interface design also has a considerable number of drawbacks. In fact, such an approach already exists and it represents the traditional manner in which interfaces for complex systems have been designed. The philosophy has been referred to as the single-sensor-single-indicator (SSSI) approach (Goodstein, 1981). Basically, it consists of displaying all of the elemental data that are directly available from sensors. Anything and everything that can be directly measured has a single display element associated with it. There are many disadvantages to such an approach (cf. Goodstein, 1981; Woods, in press), some of which are identical to those associated with rote instruments.

The most important drawback is related to controllability. In order to consistently deal with the entire range of domain demands (particularly fault situations), operators need comprehensive information regarding system state. But with the SSSI approach, only data that are directly obtained from sensors are displayed. Thus, higher-order state information which cannot be directly measured, but which is nevertheless needed to cope with many fault situations, may not be made available to operators. In fault situations, it is generally not possible to derive the higher-order properties from the elemental data, and so operators may not have all of the information that is required to consistently control the system under these circumstances.

The disadvantages of the SSSI approach are not limited to fault scenarios. Even under normal operating conditions, SSSI interfaces put an excessive burden on operators. In these situations, it is, in principle at least, possible to recover the higher-order, goal relevant domain properties from the elemental data represented in the interface. The problem, of course, is that considerable effort and knowledge are required of the operator to carry out this derivation.

Not only are the higher-order properties not directly displayed, but in addition, the relationships between the various elemental display elements are also usually not represented in the interface.

Finally, there is also the issue of information pickup. Just because the information is in the interface does not mean that the operator can find it easily (Woods, in press). In the SSSI approach, the form in which information is usually presented (e.g., similar looking analog meters, or digital numerical displays) is not very compatible with the capabilities of the human perceptual system, thereby hindering the process of information extraction. Each instrument tends to be presented individually, and there is virtually no integration or nesting of display elements. This makes it difficult for operators to perceive the state of the system, even if all of the requisite information is in the interface¹.

It is important to note that the SSSI philosophy can also be instantiated with computer-based displays, so the problems associated with rote instruments are not necessarily avoided by merely switching from hard-wired meters to computer-based displays (Woods, in press). A fundamentally different approach to interface design is required.

The disadvantages of a rote instrument approach to interface design become intuitively apparent through a simple thought experiment. What if one were to take a simple every-day task, and develop a single-sensor-single-indicator interface for it? How difficult would it be to perform the task? Take the case of tying your shoes for instance. Imagine that, instead of being able to grasp the laces with your hands, you have to issue commands to a pair of robotic arms via a command language on a standard typewriter keyboard. (Recall that in complex systems, the process cannot be directly manipulated nor observed). In addition, imagine that you cannot directly view the center of activity. Instead, you have 15 analogue, moving coil meters that separately indicate the x, y, z coordinates of: each shoelace, each robotic arm, and the shoe. Call this the single-sensor-single-indicator approach to shoe tying. The example, illustrated in Figure 1, is obviously exaggerated to make a point, but there are in fact many parallels between the single-sensor-single-indicator approach to shoe tying and traditional approaches to interface design for complex systems. If the

¹The problem of information pickup is compounded by the fact that there are so many data available, only a small subset of which is relevant for any given context. This leads to what Woods (1986; in press) refers to as the significance of data problem. The operator must determine which data are relevant for the current context. The problem of context sensitivity will not be discussed in this paper, but the interested reader is referred to Woods (1984) for a discussion of techniques for directing the operator's attention to the most relevant part of the display, to Woods (in press) and Mitchell and Miller (1986) for discussions of the problem of context sensitivity, and to Mitchell and Saisi (1987) and Woods and Elias (1988) for implementations of context sensitive interfaces. These efforts complement the ideas described in the present paper.

task seems difficult with this interface, imagine the additional problems that crop up if the "operator" had to deal with faults! How easy would it be to "see" that one's shoelace had broken? Would it be possible to feed the broken lace through the eyelet to even out the length of the lace ends, thereby making it possible to perform the task? Needless to say, this approach makes the task of shoe tying enormously more difficult than the everyday method. The point is that this conclusion is just as valid for complex work domains: current interfaces make the operator's job much more difficult than it need be.

Of course, we do not mean to imply that the difficulty associated with control of these complex systems can be reduced to that of tying one's shoes. This is clearly not possible. Controlling a power plant, for example, is a very difficult business. The point of the example is that the interface should not contribute to the difficulty of the task (cf. Flach and Vicente, 1989; Newman, 1966). Clearly, there is much room for improvement.

Figure 1. The single-sensor-single-indicator, or rote instrument, approach to shoe tying. The example illustrates the need for a smart instrument approach to interface design.

Smart Instruments

The alternative to rote instruments is what Runeson (1977) refers to as smart instruments. These are specialized on a particular type of task in a particular type of situation. Unlike rote instruments, they cannot be used for a myriad of purposes. Their great advantage, however, is that they "capitalize on the peculiarities of the situation and the task" (Runeson, 1977, p. 174). In other words, smart instruments exploit constraints that exist in a given setting to measure higher-order properties directly. The example that Runeson gives is that of a polar planimeter, a device that directly measures the area of any two-dimensional figure, regardless of its shape. The device does not use length at all to arrive at its measurement of area. No calculation, no inferences, and no knowledge of rules are required for its operation.

Runeson (1977) suggests that perception consists of smart mechanisms that directly measure complex variables. Some of the advantages of having a smart perceptual system, rather than a rote one, have already been mentioned above. There are also evolutionary reasons that give intuitive support to the idea. Being able to directly pickup goal-relevant properties that are relevant to survival would seem to be more useful than only being able to measure fundamental dimensions (e.g., Mass, Length, and Time) and then using these to painstakingly derive the higher-order properties of interest.

The idea of perception as a set of smart mechanisms is consistent with Gibson's (1966, 1979) ecological approach to perception (see also Lombardo, 1987; Michaels and Carello, 1981; Reed, 1988; Reed and Jones, 1982; Shaw, Turvey, and Mace, 1982; Turvey and Shaw, 1979; Turvey, Shaw, Reed, and Mace, 1981). From this view, organisms are able to directly perceive the affordances of the environment - the possibilities that it offers them - through a process of direct attunement, without any need for mediating inferential processes. In other words, direct perception is made possible by relying on smart perceptual mechanisms. What makes this possible is the information about the affordances of the environment. (Note that the meaning of information, as used in this context, does not imply a symbol system, cf. Turvey and Kugler, 1984). The lawful, goal-relevant relationships between the perceptual array and the organism-environment interaction provide information about affordances. The trademark of a smart instrument is the exploitation of lawful relationships - in this case, the detection of affordances in the environment.

The advantages of smart instruments suggest that the approach may be a useful analogy for interface design. As mentioned earlier, because the goal-relevant properties of the work environment are not directly available to the operator's perceptual systems, the interface must serve as the mediator

between the otherwise unobservable environment and the operator's perceptual systems. The goal of a smart approach to interface design would be to provide the information needed for controllability in a form that exploits the power of perception. Thus, there should be a nesting of invariants in the interface, thereby allowing the operator to directly perceive the affordances of the work domain at any level of abstraction. As is the case with the natural environment, we would like to create a flexible interface that does not constrain operators to a certain level, but instead, allows them to observe the system at the level which is most appropriate for the given context (Flach and Vicente, 1989; Rasmussen and Vicente, in press; Vicente and Rasmussen, 1989). Such a "smart" approach to interface design would overcome many of the problems associated with the SSSI approach.

The smart instrument analogy suggests several important questions that must be addressed if the interface is to serve as an effective mediator:

1. Content - What are the goal-relevant properties of the environment that need to be measured?
2. Structure - How are these different properties related?
3. Form - What visual form should these properties take?

These questions, which map onto the three dimensions of content, structure, and form, define the core of the interface design problem. Because these questions are so central to the concerns of this paper, it is important to delve into them in more detail. The first question - what properties need to be measured? - demands an analysis of the objectives of the work domain and the various means for control that are available for accomplishing those objectives. Such an analysis will determine which properties of the work domain are relevant to the operator. These are the properties that should be displayed in the interface. As already mentioned, SSSI interfaces may fail to provide all of the goal-relevant information that the operator needs to cope with abnormal events. Clearly, a more comprehensive approach is required.

Ecological psychology has developed the notion of an affordance to approach similar problems (Gibson, 1979). What is needed for interface design is something equivalent to an analysis of the affordances of the work domain. That is, designers need to analyze the effectivities of the work domain (i.e., the available control capabilities) so that they can determine what the work domain affords. This should enable the designer to construct a set of smart display elements that measure the goal-relevant variables (cf. Woods, 1986), thereby providing operators with the action-scaled information (Warren, 1985) they need to control the system over the entire range of domain demands.

Conducting a thorough analysis of a work domain's affordances will overcome the controllability problems associated with the SSSI approach.

The third question - what form should the properties take? - requires a deep understanding of the operator's perceptual, action, and cognitive capabilities. To the extent that these are limited, they represent a set of constraints on interface design. As mentioned above, information in SSSI interfaces is typically presented in a form that is very difficult for operators to pickup. What is required is some way of conveying the affordances of the domain in a way that makes them directly available to the operator's perceptual systems.

Again, ecological psychology has developed a relevant theoretical tool that addresses this issue. Ecological optics is the analysis of the structure in the ambient optic array, and in particular, how that structure is specific to the affordances of the environment (Gibson, 1961). Extrapolating to the present concerns, the task of the interface designer is actually one of inverse ecological optics. That is, given an understanding of the human perceptual system, the designer should use computer technology to make the previously identified affordances available to the organism in a form that "vision is ready to pickup" (Gibson, 1979, p. 260) . Providing information which specifies the affordances of the work domain will allow for direct perception.

Finally, answering the second question - how are the different properties related? - will provide a basis for organizing the information in the interface. This is equivalent to determining how the various affordances of the ecosystem are related. SSSI interfaces group display elements into functional groups, but do not reveal much of the higher-order domain structure. Revealing the semantic structure of the work domain in the visual structure of the interface should make it easier for operators to apprehend the displayed information, since this is the very information they need for control. Unlike the previous two questions, however, ecological psychology has not developed a tool for addressing this problem. In the next section, a framework that is suited for this problem will be described.

Summary

Understanding how human perception functions has led to some important insights for interface design. In particular, the discussion suggests that a smart instrument approach to interface design is needed to supplant the rote instrument (single-sensor-single-indicator) approach which has dominated interface design. In addition, several conceptual tools (i.e., affordances and ecological optics) which can perhaps be meaningfully adapted to the problem of interface design have been uncovered. This is not to say that ecological psychology has ready-made answers for these problems. This is certainly not

the case. The point is that the concepts of affordances and ecological optics can be used to talk about parallel issues associated with interface content and form, respectively. However, it is not obvious how, or whether, these concepts can be used in a prescriptive manner as is required for design.

The one thing which seems to be missing is a format for structuring the affordances of a work domain. This is a problem that ecological psychology has yet to approach in a comprehensive manner. Affordances have typically been studied as single isolated entities (e.g., pass-throughable, sit-onable, graspable). But in any realistic situation, a large number of inter-related affordances are available to the active organism. It would be useful to examine the structure that exists both between and within affordances, since this structure may convey important goal-relevant information.

Rasmussen (1985a, 1986) has developed a conceptual tool, called a means-end or abstraction hierarchy, to approach a similar set of problems. In the following section, the concept of a means-end hierarchy will be described as it relates to the problems of ecological psychology. While it was originally developed to describe the functional structure of complex work domains, we will show how it can also be used to bring out the relationships between the various affordances of the natural environment.

THE MEANS-END HIERARCHY: A FRAMEWORK FOR COPING WITH COMPLEXITY

Figure 2 presents a number of the examples that Gibson (1979) used to illustrate the concept of an affordance. While all of these examples share the common property of being a possibility for action afforded to the organism by the environment, there also seem to be differences between some of the examples as well. Intuitively, some affordances in Figure 2 seem to be more similar than others. For example, Lifting seems to be more similar to Cutting than to Nurturing or Survival. The question we would like to pose in this section is: How are these different affordances related? The goal will be to describe a framework for organizing affordances into a coherent structure.

Developing A Means-End Hierarchy of Affordances

Where is an appropriate place to start in uncovering the structure among and between the affordances listed in Figure 2? Hierarchical structures have been found to be useful in representing complex systems in a variety of different applications (e.g., Allen and Starr, 1982; Korf, 1987; Mesarovic, Macko, and Takahara, 1970; Pattee, 1972; Simon, 1981). Thus, a good initial guess would be to identify different levels of description for ordering affordances. In fact,

Gibson (1979, p. 137) briefly alludes to a hierarchy of affordances, but does not pursue the idea any further.

Figure 2. A sampling of affordances of the natural environment, taken from Gibson (1979). A question that has yet to be effectively addressed is: How are these different affordances related?

An immediate problem that arises is that hierarchies come in many different flavors (cf. Mesarovic et al., 1970). One of the key characteristics distinguishing different types of hierarchies is the relation between levels. Many different relations are possible (e.g., spatial scale, temporal scale, authority, flow of information, etc.). Which is most appropriate for the current context? Given the view of organisms as goal-directed systems, a means-end relation seems to be appropriate. What would a "means-end hierarchy" look like, and more importantly, how useful is it in providing a framework for describing the set of affordances illustrated in Figure 2?

Figure 3 illustrates some of the relations that can be derived through a means-end analysis of affordances. Three levels are represented, each representing three different questions that are crucial to any goal-oriented system: WHY? WHAT? and HOW? The three levels are not absolute in any

sense, but rather can "slide" up or down (this point will be described in more detail below). The level at which the observer enters the system defines the WHAT level. Given this choice, the level above specifies WHY and the level below specifies HOW. A concrete illustration will help in providing a more intuitive understanding of this structure.

Figure 3. Coping with complexity: exploiting a hierarchy of affordances. The figure shows how there are constraints between affordances, and how the means-end hierarchy can be used to bring out this structure.

To begin with, take the affordance of Locomotion in Figure 3 as an example. If, for some reason, an organism was interested in locomoting from point A to point B, then it would be important to know what means are immediately available for achieving this goal. As shown in Figure 3, these means are specified at the level below (i.e., HOW is it possible to locomote?). Possible means for locomoting include Crawling, Flying, Walking, and Climbing. Similarly, relevant means for an organism who is interested in playing include Throwing, Flying, and Climbing.

Returning to the function of Locomotion, possible justifications or reasons for wanting to locomote are specified at the level above (i.e., WHY would one want to locomote?). Referring to Figure 3, two reasons for wanting to locomote are shown: Privacy (e.g., to escape from a crowd) or Cooperation (e.g., if the partner is on the other side of the pool). Similarly, the act of Playing can be undertaken to serve the goal of Cooperation. In a sense, these higher order purposes provide the context for the current affordances of interest.

Several points about the structure shown in Figure 3 are worthwhile pointing out. First, as mentioned before, the three levels defined by the questions, WHY, WHAT, and HOW are relative. For instance, if the observer entered the system at the level of Privacy, then that level would be the WHAT level, and the levels above and below would be those of WHY and HOW, respectively. Second, it is also important to note that there can be a many-to-many mapping between affordances at various levels. Thus, a certain affordance (e.g., Locomotion) can be fulfilled by a number of different alternatives. Also, a single affordance (e.g., Swimming) can be related to more than one higher-order affordance. A third

feature of the hierarchy is that it represents the constraints that exist on achieving goals. For example, Throwing is not a meaningful means for achieving the goal of Locomotion. The most important advantage of adopting such a format to structure the affordances of an ecosystem is that it provides a mechanism for coping with complexity (Rasmussen, 1985a). The hierarchy in Figure 4, consisting of all of the affordances in Figure 2 along with a few others, illustrates how higher-order affordances can help people to cope with the complexities inherent in the natural environment (the constraints between levels, as represented by the links in Figure 3, are omitted for clarity). To avoid misunderstanding, it is important to note that the hierarchical structure in Figure 4 does not imply that higher level affordances are constructed by integrating information from lower levels. Consistent with Gibson's (1979) position, any of the levels can be directly perceived.

For any given situation, perception of lower-level affordances may be constrained by the perception of higher-order, value-related affordances. For example, whether a surface is Swim-overable will be directly relevant if the current context affords Locomotion, but not if it affords Drinking. The means-end relation thereby allows people to identify the goal-relevant possibilities with respect to relations between levels. "The environment is a continuum of information sources, and to cope with this great variety, the [organism] must be able to access this information in manageable higher-level structures. This means that information pickup is an active process that may be organized top-down in a means-end hierarchy" (Rasmussen, 1986, p. 79).

The important implication of this discussion is that perception of the situational context in terms of higher-order affordances can help organisms deal with the complexity of the environment by constraining the number of meaningful lower-level affordances which are relevant, given the current context. A means-end hierarchy of affordances reveals emergent goal-relevant information in the relations between and among affordances that is not contained at the level of individual affordances.

Value Properties: Purpose, Goal			
Survival	Pleasure	Altruism	
Priorities: Abstract Function			
Reward	Danger	Nutrition	Manufacture
Cooperation	Nurturing	Copulation	Privacy
Comfort	Pain		
Context: General Function			
Communicating	Warmth	Drinking	Eating
Washing	Bathing	Injury	Support
Fighting	Shelter	Aiding	Punishment
Locomotion			
Movement: Physical Process			
Sit-on	Bump-into	Fall-off	Get-underneath
Climb-on	Sink-into	Swim-over	Walk-on
Stand-on	Grasp-able	Barrier	Obstacle
Breathing	Pouring	Cutting	Lifting
Throwing	Piercing	Carrying	
Objects and Background: Physical Form			
Layouts	Objects	Surfaces	Substances

Figure 4. Affordances structured within a means-end hierarchy.

Generic Properties

In this section, the generic properties of Rasmussen's means-end hierarchy will be described in abstract terms. In addition to discussing the nature of the hierarchy in more detail, the discussion will also draw connections to an important general class of problems.

First of all, it is important to be precise about the way in which the concept of hierarchy is being used. The definition of hierarchy that has been adopted here is based on the treatment provided by Mesarovic et al. (1970), which, incidentally, differs from the definition used by Turvey, Shaw, and Mace (1978). The means-end hierarchy belongs to the class of stratified hierarchies described by Mesarovic et al. (1970, pp. 37-43), the properties of which are listed below.

1. Each stratum, or level, of the hierarchy deals with the very same system, the only difference being that different strata provide different descriptions, or different models for observing the system.
2. Each stratum has its own unique set of terms, concepts, and principles.
3. The selection of strata for describing a particular system depends on the observer, and his knowledge and interest in the control of the system. For

many systems, however, there may be some strata which appear to be natural or inherent.

4. The requirements for proper system functioning at any level appear as constraints on the meaningful operation of lower levels, while the evolution of the state of the system is specified by the effect of the lower levels on the higher levels.
5. Understanding of the system increases by crossing levels: by moving up the hierarchy, one obtains a deeper understanding of system significance with regard to the goals that are to be achieved, while in moving down the hierarchy, one obtains a more detailed explanation of the system's functioning in terms of how those goals can be carried out.

Armed with this description of hierarchy, let us now proceed to the particulars of the means-end hierarchy and the problems it is meant to address.

An omnipresent problem associated with complex systems, whether it be a complex work domain or an ecosystem, is what is generically known as the degrees of freedom problem (e.g., Bernstein, 1967; Saltzman and Munhall, 1989; Turvey, 1977; Turvey et al., 1978). In the case of technical systems, the problem can be posed as follows: How can an operator with limited resources decide what actions will lead to the satisfaction of system goals, if the system consists of so many components with many interactions between them? Clearly, for resource-bounded agents to be able to deal effectively with such complexity, they must somehow reduce the dimensionality of their description of the system of interest.

The means-end hierarchy provides a mechanism for coping with such complexity. Going up the levels of the hierarchy, there is a selective loss of detail, and thus problems become more manageable. A formal analysis of the computational advantages that result from such a hierarchical structure has been conducted by Korf (1987). However, the loss of detail is not the entire story. Not any type of hierarchy will do. The critical aspect is the means-end relation between levels. This relation allows goal-directed organisms to traverse the hierarchy in a way that will make it possible for them to solve complex problems. For instance, the organism can observe the system at a higher level and then decide which part of the system is of greatest relevance to its immediate goals. By moving down to the next level, the organism can then consider only the factors that are functionally connected to that part of the system. In other words, the constraints revealed in the structure of the hierarchy allow organisms to focus their attention increasingly on the most appropriate (i.e., goal-relevant) component of the system by "navigating around" through the various levels of the hierarchy. Due to the complexity of the system, exhibiting proficient goal-directed behavior would be very difficult -

perhaps impossible - if the organism had to observe the system at the finest grain. The level of detail would simply overwhelm the organism's limited capabilities. While there are cases where such detail is necessary, they represent a minority.

The means-end hierarchy, then, represents the functional structure of a system (i.e., the set of constraints on achieving the system's objectives). It has a simple yet intimate connection to the ecological concepts of affordances and effectivities: the means available for control are effectivities, and the ends to be achieved are affordances. Thus, the hierarchy is intended to be a representation of the inventory of all available means related to all relevant ends and constraints, seen as boundaries around the acceptable pursuit of the system goals. It describes the available variety for coping with all relevant situations (i.e., the possibilities of the system, or requisite variety in the sense of Ashby, 1956). It is similar to the concept of an epistemic ecosystem (cf. Turvey and Shaw, 1979; Turvey et al., 1978; Shaw and Turvey, 1981) in that it is closed. When used to represent a work domain, the hierarchy defines the functional "landscape" (i.e., topology) through which agents will "navigate" during their work (see Rasmussen, 1986, p. 119 for an example from the domain of electronic troubleshooting). Individual control trajectories can be mapped onto the space, but no control trajectory, or even controlling agent is explicitly represented in the hierarchy.

Summary

In this section, the means-end hierarchy has been introduced as a generic mechanism that can be used to describe complex, goal-oriented systems with many degrees of freedom. The framework is relevant to interface design because it provides a way to structure the properties that are to be represented. However, we have tried to show that the very same structure can also be meaningfully adopted to organize the affordances of an ecosystem. The constraints represented by the mappings between levels of the means-end hierarchy show how it is that "occasions [i.e, situational contexts] individuate affordances" (Turvey et al., 1981, p. 299). This development represents a contribution to the conceptual framework of ecological psychology.

This is not to say that the means-end hierarchy is a mature formalism, although it has been successfully applied to a variety of domains in cognitive engineering (e.g., Rasmussen, 1988b; Woods and Hollnagel, 1987). It follows from this caveat that the affordance structure illustrated in Figure 4 is not meant to be read as being definitive or "correct". The intention has merely been to illustrate the properties of the means-end hierarchy, and to point out its relevance to the questions of ecological psychology. We hope that future work

will deepen and strengthen the relationships that have been pointed out. Providing a detailed description of the means-end hierarchy within the conceptual framework of ecological psychology, as we have done here, is a first step in this direction.

Other Applications to Ecological Psychology

There are several potential applications or extensions of the means-end hierarchy that are relevant to ecological psychologists. For one, it would be useful to try to formalize the framework. A plausible move in this direction would be to apply dynamics (Abraham and Shaw, 1982, 1983, 1984, 1988) to quantify the topological properties of the different levels in the means-end hierarchy. It would also be beneficial to understand the relationship between the means-end hierarchy and the formalization of an ecosystem as a coalition, as described by Shaw and Turvey (1981). Such comparative analyses could lead to the integration of the conceptualizations, or at least, to a better understanding of how they differ.

It is also likely that the means-end hierarchy can be applied to motor control. Like an affordance, the concept of a coordinative structure is described functionally in a goal-specific manner (cf. Greene, 1982; Kugler and Turvey, 1987; Saltzman and Munhall, 1989; Turvey, 1977; Turvey et al., 1978). Given the insights that were derived from the means-end analysis of affordances, it would be worthwhile to conduct an analogous investigation to understand how various coordinative structures are related and constrained. The process of organizing coordinative structures into a means-end hierarchy should reveal goal-relevant information that is not apparent when the individual coordinative structures are viewed in isolation. For instance, the mappings between levels would specify which lower level coordinative structures could be selected to implement the functional requirements of the higher level units.

Now that a framework for representing the structure of a complex work domain has been introduced, it is possible to return to the problem of interface design.

ECOLOGICAL INTERFACE DESIGN

It is time to assemble the insights that were gathered from the discussions on smart instruments and the means-end hierarchy and bring them to bear on the problem of interface design for complex systems. The outcome of this endeavor is Ecological Interface Design (EID), a theoretical framework for interface design for complex work domains (Rasmussen and Vicente, 1989; Vicente and Rasmussen, 1989).

Different Types of Constraints in Complex Work Domains

The discussion of perceptual instruments suggests that a smart instrument approach to interface design would be appropriate. As mentioned earlier, the key to smart instruments is that they exploit constraints. Fortunately, in the case of complex work domains, there are several different types of constraints that can be exploited by a "smart" interface. First, there are global constraints, which are the purposes for which the system was designed. For example, in the case of a power plant, the global invariants are to produce power and to do so safely. In addition, there may also be holonomic constraints, or laws, which describe the functioning of the system. To take the example of the power plant again, the behavior of the system is governed by the laws of thermodynamics, which are related to the conservation of energy and mass (cf. Beltracchi, 1989). Finally, there are also nonholonomic constraints that can be exploited. These can be thought of as the boundary conditions that are set up by the initial design of the system (e.g., the number and type of components, the way in which they are configured, and so on). It is evident, then, that there is a great deal of structure in these systems, making a smart instrument approach feasible.

The means-end hierarchy was suggested earlier as a way to represent the structure of a complex system. In fact, the hierarchy can represent all of the different types of constraints mentioned in the previous paragraph (i.e., global, holonomic, and nonholonomic). Furthermore, it can also show how the constraints are related between levels of the hierarchy. This suggests that the means-end hierarchy provides a fundamental basis for interface design, specifying, not only which set of properties should be included in the interface (i.e., content), but also the relationships between those properties (i.e., structure). The third and final step would be to present these properties in a form that is easy to perceive.

Revealing the Domain Constraints in the Interface

In deciding on a visual form, it is important to keep in mind that the information in the interface can be interpreted in qualitatively different ways (Rasmussen, 1983). Although the goal is to allow operators to control the system by relying on perception and action as much as possible, as mentioned earlier, there will inevitably be times where problem solving activities will be required to deal with domain demands. Therefore, it is critical that the very same interface provide information to allow operators to effectively control the

system in either processing mode (cf. Bødker, 1989; Vicente and Rasmussen, 1989)¹.

Summarizing, the goal of EID is twofold: first, to reveal the invariants of the work domain through the interface in such a way as to take advantage of the powerful capabilities of perception and action; and second, to provide the appropriate computer support for the comparatively more laborious process of problem solving.

The Principles of EID

The EID framework consists of three prescriptive design principles for building interfaces that satisfy this goal (see Vicente and Rasmussen, 1989 for the theoretical development of these principles within a cognitive engineering framework).

1. Use the means-end hierarchy for representing the work domain as a hierarchy of nested affordances. This representation specifies the content and structure of the interface, and provides a normative basis for problem solving.
2. Map the semantics of the ecology onto the geometry of the interface in order to reveal the affordances of the work domain in a way that exploits direct perception. The goal here is to "make visible the invisible".
3. To support interaction via the perception-action cycle, operators should be able to directly act on the display, and the structure of the displayed information should be isomorphic to the part-whole structure of movements.

Although these principles may seem abstract, this is necessary if the approach is to be a general one. The intent is to get designers to ask useful questions, not to provide a "design cookbook" that can be followed by rote. The latter approach is not feasible since the demands, and therefore the detailed design solutions, for each domain will be unique. Later in the paper, we will show how these abstract principles can be instantiated to produce a concrete design product.

Before describing an application, however, several assorted topics which are pertinent to EID will be discussed. The remainder of this section will be concerned with: elaborating on the role of action in interface design; discussing the flexibility of control made possible by EID; the relationship between

¹This requirement implies that there may be tradeoffs involved. For instance, it may be possible to design an interface that would exploit perception more effectively, but doing so could perhaps inhibit effective problem solving. This would not be an acceptable design, since it is in situations where the operator needs to "go beyond the information given" that the risk for disaster and the need for help from the computer system both reach a peak. Thus, it would be preferable to develop an interface that provides better support for problem solving, even if it means that one has to sacrifice the ease with which it supports perception. Future research should be directed at determining whether or not such a tradeoff exists, and if so, how to strike an effective balance between these two conflicting goals.

interface design and training; and finally, the connection between direct perception and EID.

What About Action?

The comparative lack of emphasis that EID places on action may be surprising, since perception and action are duals of each other. This asymmetry will become particularly apparent in the example discussed in the following section. There are at least two reasons for not giving action a more prominent place. First, comparatively little is known about manual control of systems with very slow dynamics. The studies that have been performed suggest that, with slow systems, the importance of the coordinative aspects of movement tends to diminish and is instead replaced by a greater emphasis on the strategic aspects associated with discrete control (e.g., Crossman and Cooke, 1974). Second, issues of action within the context of process control naturally encompass the engineering domain of control systems design, since the action capabilities of the system are a joint function of the properties of human manual control and those of existing automatic control technologies. The interaction between these two areas has yet to be addressed in a comprehensive manner.

It is worthwhile considering how control systems design can be integrated with the ideas about interface design expressed in this paper. An obvious suggestion is that it would be useful to have perception and action occur at the same level. The EID interface in the example below allows for direct perception of any level, but control is possible only at the elemental level of physical components. In the proposed scheme, control of the system could proceed at any level of the means-end hierarchy. One can envisage a hierarchy of display screens corresponding to the various levels of the means-end hierarchy (cf. Goodstein, 1983), and within each screen, perception and action would be directly coupled at the same level of abstraction. To implement this idea, one would need to construct a hierarchical control algorithm that would deal with the degrees of freedom problem associated with implementing a command at a high level of abstraction. Furthermore, it would also be necessary to design the automatic controllers so that functions at the same level of abstraction are decoupled from each other, thereby allowing operators to change one function without creating an unwanted change in another. This is a tall order since the many-to-many mapping between levels that typically exists in these systems makes it likely that some higher-order functions will be coupled at lower levels of abstraction.

These speculative comments lead to the perhaps surprising conclusion that the design of automatic controllers for complex, technological systems can perhaps benefit from psychological considerations, particularly the importance

of directly coupling perception and action. The strategy of bringing psychological issues to bear on control systems design was adopted many years ago, and with great success, within the limited context of vehicular control systems (cf. Flach, in press a; Wickens, 1986). Although the exact nature of the issues are different from those encountered in vehicular control, the same general strategy should be considered within the context of complex work domains. Interestingly, this idea is also consistent with the view of the design of the (work) environment as the design of affordances (cf. Gibson, 1982a; Warren, 1985).

Flexibility of Control

The EID approach to interface design has an interesting property with regard to flexibility of control. Because the interface content is based on the means-end hierarchy, the operator is free to choose whatever means are available to satisfy any given function. This contrasts with the traditional human factors approach to design which is typically based on behavior rather than structure. The classic approach would be to conduct a task analysis to identify a single sequence of overt behaviors for performing each task (e.g., Meister, 1985). Following this philosophy, the design would be optimized for that particular way of performing the task, but it would not necessarily support other control strategies. On the contrary, it may even impede other strategies.

As mentioned earlier, the means-end hierarchy does not have any control trajectories embedded in it. Thus, operators are free to choose the strategy they prefer, as long as it is consistent with system goals. An interface based on the principles of EID will attempt to provide an informational basis for behavior by representing the set of constraints relevant to the satisfaction of system goals, but it will also allow the operator to work within his subjective preferences. This will result in a naturalness of control that is not available from interfaces that are based on a procedural task analysis. In addition, such a design approach has the further advantage of providing the flexibility that is required to adapt to contingencies that cannot be anticipated by the designer (cf. Rasmussen, 1985a). This is a critical requirement in the control of complex work domains.

Learning and Training

In a top-down approach to system design, issues of training and learning go hand in hand. Therefore, it is important that the respective roles of training and interface design be clearly specified (cf. Vicente, 1988). The first question which must be addressed is: What is the population that the interface is being designed for? In our case, the design is intended for experienced operators of systems who have gone through rigorous training programs and who use the

interface in their daily work activities. This simplifies things to a certain extent since problems associated with periodic use, or negative transfer between systems are not of great concern.

Another important characteristic of complex work domains that bears on learning is that there is often a normative basis for system functioning (e.g., the laws of thermodynamics in the control of a power plant). If the operators are to exhibit effective control, they must work within the constraints that govern the behavior of the system they are controlling. Thus, in addition to serving as an effective mediator, the interface can also play an instructional role by educating the operator's attention. From this perspective, the interface should allow the operator to induce effective control strategies through experience. EID attempts to facilitate induction by making the goal-relevant constraints visible. Through the learning phase of active, goal-directed exploration, the operator should gradually pickup on these constraints, and develop appropriate control strategies (see the example in the next section). Ideally, the interface would allow the operator to acquire a complete, veridical understanding of system behavior, merely through extensive practice with controlling the system.

In real systems, however, interface design alone will not do the job. Because of the sources of complexity mentioned earlier, the operator's attention will also have to be educated through instruction, not just by experience. This will be especially critical in fault situations, since the operator will not usually have prior experience with the particular fault he is faced with. Thus, training complements interface design by providing operators with a more general understanding of system functioning that is not tied to the limited set of situations they have experienced (Vicente, 1988).

Mediated or Direct?

In previous discussions of EID (Rasmussen and Vicente, 1989; Vicente and Rasmussen, 1989), it has been argued that by following the principles outlined above, it will be possible to design interfaces that allow operators to rely on direct perception (see also Flach and Vicente, 1989). However, earlier in the paper, the interface was described as a mediator. How can perception be direct if it is mediated by the computer?

The first point that needs to be made is that perception can be more or less direct (Gibson, 1982b; Mace, 1980). As the opening quote above from Gibson (1979) indicates, the direct/indirect distinction lies on a continuum. As one moves along this continuum in the direction of increasing indirectness, the information available for perception is less rich, and thus, there will be a greater need for inference. For example, looking at objects that are very far away via a telescope is indirect in that it does not enable one to explore the

structure in the optic array, as would be possible if the object was close to the perceiver. Inference may be needed to compensate for the decrease in information availability. Apprehension via computers is even more indirect in that the computer plays the role of a medium. The primary difference between direct perception via an optic array and apprehension via an ecological interface is that light as a medium is completely transparent, but the computer is not. There will always be information specific to the computer's own properties as an object in the environment distinct from the system being controlled (cf. Gibson, 1979, pp. 281-283).

The goal of EID then, is to make the computer as transparent as possible, thereby minimizing the need for inference. Just as the nervous system and the energy medium are functionally transparent to the properties of the world that are relevant to the survival of the organism (Shaw and Bransford, 1977), so the computer interface should be as functionally transparent as possible to the properties of the work domain that are relevant for effective control (i.e, for the attainment of system objectives).

The idea is to "make visible the invisible" so as to create the phenomenological feeling in operators that they are directly controlling the internal functions of the system, not dealing with a computer intermediary (cf. Laurel, 1986). This is accomplished through direct specification. The need for inference is minimized because the interface provides an isomorphic mapping of the goal-relevant properties. The interesting point is that by achieving direct specification, the very same interface can be interpreted as a signal for perceptual-motor control, as a sign triggering an action rule, and as a symbol for problem solving (Rasmussen, 1983), thereby supporting all modes of processing that the operator will need to rely on (cf. Vicente and Rasmussen, 1989).

The idea of direct specification can be considered as an example of apprehension mediated by instruments (Gibson, 1982b). In designing direct specification instruments, irrelevant information will be left out while goal-relevant information will be enhanced. An example is perception in the dark via infra-red goggles. Certain invariants in the optical array are preserved (e.g., shape of objects), whereas others (e.g., color) are not.

An important implication associated with apprehension mediated by instruments is that: "the reality testing that accompanies the pickup of natural information is missing The invariants have already been extracted. You have to trust the original perceiver" (Gibson, 1979, p. 261). In the case of apprehension via ecological interfaces, it would be more appropriate to say: "You have to trust the designer". That is, informative perception mediated by instruments is only possible if the designer has constructed the instrument in

such a way that it conveys the intended invariants (e.g., consider the knowledge that must go into designing an electron microscope). In the case of interface design, the extraction of the invariants in the work domain is performed analytically beforehand by the designer and then built into the interface geometry. Thus, there are a fixed number of properties that the operator is receiving information about.

It follows that there are some properties which are not being conveyed. There may be some situations (e.g., a fault causing a structural change in the system) where new properties that are not represented in the existing interface may become relevant to the control of the system. As an example, if one were tying one's shoes with the SSSI interface in Figure 1 and one of the laces broke, the property of lace tension would be helpful in diagnosing the "fault". But this property was not considered by the designer as being relevant to system goals, and therefore is not available in the interface. Kugler and Lintern (1989) provide an interesting discussion based on the concepts of self-organization that could potentially address the issue of designing interfaces that cope with the emergence of novel, goal-relevant properties.

It should be pointed out, however, that adopting the means-end hierarchy as a basis for determining which properties need to be represented in an interface should allow operators to control the system under abnormal conditions. The reason for this is that, regardless of the nature of the abnormality, the set of means (effectivities) available for coping with the fault is bounded by the initial design of the system, and by the holonomic invariants (i.e., laws), both of which are represented in the means-end hierarchy. Because the holonomic invariants (e.g., the conservation of mass and energy) always hold, they provide a basis for control in abnormal situations.

Summary

In this section, the principles of EID have been described. In addition, issues associated with the role of action, flexibility in control, training and learning, and direct perception have been discussed as they pertain to EID. These are factors that need to be taken into account in a comprehensive approach to system design.

AN APPLICATION OF THE PRINCIPLES

In this section, the principles that were described above will be illustrated within the context of a thermal-hydraulic process control system. While the interface developed below will serve as the basis for experiments evaluating the EID framework, it should not be taken as a final product, but rather as an

initial prototype. It is likely that an empirical evaluation will reveal how the form of the interface can be improved.

It is important to note, however, that the content of the interface is not subject to empirical evaluation. The reason for this is that interface content can, and should, be evaluated analytically (Rasmussen, 1985b; Vicente, 1990). The role of experimentation is to determine how that information should be revealed (i.e., what form it should take). Thus, rather than illustrating a definitive design, the purposes of the example are: first, to show that the principles of EID can be instantiated for a given domain; and second, to point towards the need for more research on principle-driven design for direct specification.

The Ecology of DURESS

DURESS (DUal REservoir System Simulation) is a thermal-hydraulic process simulation, to be used as a research vehicle (Vicente, 1987). If the findings of basic research are to generalize to real world systems, it is imperative that the simulation have a minimum degree of complexity. While it does not embody the full complexity of an actual work domain, DURESS is considerably more complex than the tasks that have traditionally been adopted for basic laboratory research.

DURESS represents several of the factors that make real-world problem solving difficult (cf. Woods, 1988). First, the system is dynamic, with time lags on all of the control variables. Thus, there is a considerable delay before the effects of control actions result in a visible change in system state. In addition, effects also take some time to propagate to different parts of the system. Second, DURESS also has couplings between subsystems. In some situations, this puts a burden on the operator to cope with multiple, competing goals (see below). Third, there is also a degree of risk since the reservoirs can be damaged by heating them while they are empty. Damage can also be incurred by allowing the water to spill from the reservoirs. As a result of its complexity, DURESS poses, on a much smaller scale, many of the demands that designers of high-tech systems are likely to face. It serves, therefore, as a challenging context for testing the principles of EID.

The physical structure of DURESS is illustrated in Figure 5. The system consists of two redundant feedwater streams that can be configured to supply water to two reservoirs, or vats. The operator's goals are to keep each of the reservoirs at a prescribed temperature (40 °C and 20 °C), and to maintain enough water in each reservoir to satisfy each of the current externally determined demand flow rates. To do this, the operator has control over six valves, two pumps, and two heaters. The diagram representation format

illustrated in Figure 5 has often been adopted as a basis for interface design. Our contention is that such an interface format is not very effective in supporting operators in coping with the entire range of control demands that they are likely to face (e.g. normal operations, startup, shutdown, and disturbances). The reason for this is that the interface in Figure 5 represents the relational structure at the level of physical anatomy, whereas process operation has to be planned at the level of functional relations. This point perhaps can be made more obvious by developing an ecological interface for DURESS.

Figure 5.A physical diagram interface for DURESS.

Describing the Affordances of the Ecology

The first step in designing an ecological interface is to represent the affordances of the work domain. For this purpose, EID adopts the means-end hierarchy. As mentioned before, this hierarchical domain representation specifies the content and structure of the interface. While not explicitly described below, the various levels of the means-end hierarchy can also be decomposed along a part-whole hierarchy. The part-whole dimension allows operators to view a given level at different grains of resolution (see the example of the two feedwater streams described at the level of Physical Function below).

Figure 6 illustrates the means-end hierarchy for one reservoir. It should be noted that the system's dynamic equations provide a formal description of the qualitative structure represented in the figure. At the top level of Functional Purpose, the higher level objectives of the domain are specified: to keep the water in each temperature at criterion, and to keep enough water in each reservoir to satisfy the respective output demand. While not explicitly represented, there are also constraints on how these objectives are to be met (e.g. avoid spilling or boiling water, do not damage any components, etc.).

Figure 6. A means-end hierarchy representation of DURESS showing the relationships within levels of the hierarchy.

It is possible to describe the levels of the means-end hierarchy in ecological terms. For instance, the level of Functional Purpose represents the higher-order affordances of the domain. Each of DURESS' reservoirs affords the supply of water at a certain temperature. The level below will specify the effectivities (i.e., the means) that are available for control. This relationship (means $\rightarrow$ end, or affordance $\rightarrow$ effectivity) holds between each pair of adjacent levels in the hierarchy.

At the second level, Abstract Function, the causal structure of the system is described according to first principles; in this case, conservation of mass and

energy. Thus, at this level DURESS is represented in terms of its mass and energy topology. For the mass balance, there are three variables: the rate at which mass is flowing into the reservoir (source), the rate at which mass is flowing out of the reservoir (sink), and the current mass inventory, as reflected by volume. The energy topology is similar, the only difference being that there are two energy sources, the incoming water and the heater.

At the level of Generalized Function, DURESS can be described in terms of several standard heating and liquid flow functions that instantiate the mass and energy topology specified above. Thus, in the energy flow, there is a heating source, a cooling source, a means for removing heat, and a means for storing heat. In the mass flow, there is a means for adding water, a means for removing water, and a means for storing water. At this level, we see that the energy and mass flows are actually coupled: the source of water is also a cooling source; the means for storing water is also a means of storing energy; and the means of removing water is also a means of removing energy.

At the level of Physical Function, the system is described in terms of the components that realize the functions specified at the level above. These are the variables the operator has control over. Due to space restrictions, the feedwater streams are represented as a whole in Figure 6, but a more comprehensive representation would decompose each stream into its constituent components (i.e., three valves and one pump) along a part-whole dimension. While not explicitly shown in Figure 6, at this level we also note that the two reservoirs are coupled by the redundant feedwater streams. Either feedwater stream can be configured to feed either reservoir.

The level of Physical Form is not shown in Figure 6 because it is similar to the level of Physical Function. The difference is that whereas Physical Function specifies the settings and causal relationships between components, Physical Form specifies the spatial layout and appearance of those very same components.

Whereas Figure 6 illustrates the relationships within levels of the means-end hierarchy representation, Figure 7 shows the relationships between levels. These mappings between levels reveal some important structural system properties. There are three properties, in particular, that are worth noting.

Figure 7. The mappings between the levels of the means-end hierarchy for DURESS.

First, the mappings between Functional Purpose and Abstract Function indicate that the temperature goal is connected to both the energy and mass balances. The reason for this is that temperature is defined as energy per unit mass. In contrast, the volume goal is only connected to the mass balance. In practical terms, this implies an inter-goal constraint (Woods and Hollnagel, 1987). Making a change in the mass flow topology to affect volume may have the unintended side effect of changing temperature. Given the goals of the system, this is an unavoidable fact. It results from the fact that one of the goal variables (temperature) has the other goal variable (volume) as one of the factors that enters into its definition.

A second property of interest results, not from a fundamental consideration, but from the particular physical configuration that was selected for the system. In particular, there is a many-to-one mapping between Abstract Function and Generalized Function. There are several instances where a single subsystem is a means for controlling both the mass and energy balances. For instance, each feedwater stream serves as both a source of mass and a source of energy. Thus, there is a structural coupling between the mass and energy configurations. Note that these couplings would not exist if separate components were implemented to instantiate the mass and energy balances.

The third structural property is similar. Figure 7 indicates that there is a many-to-one mapping between Physical Function and Generalized Function.

This is because either feedwater stream can be used to supply either vat. In other words, there is an indirect coupling between vats via the feedwater streams. This means that changing the settings of the valves or pump in one stream to control the amount of water flowing into one of the vats could actually inadvertently affect the flowrate to the second reservoir. Of course, this problem could be eliminated if each feedwater stream served only one vat, but then the capability to deal with faults would be removed. The benefit of having two redundant streams is that, in the case of a fault in one stream (e.g., a blocked valve or pump), the system can be reconfigured so that the other stream supplies water to both of the vats. In this way, the system will still be able to function as desired.

It is important to understand how these last two features of the system are related to the first one discussed above. The three are similar in terms of their practical implications: changing the setting of one component may have unintended side effects. However, the properties differ in terms of their origin. The first property is unavoidable (given the current Functional Purpose) because it has a fundamental basis, whereas the other two are, in principle at least, avoidable because they result from specific design decisions.

The preceding discussion reveals how the means-end hierarchy can provide some important insights into the structure of DURESS. Some other properties of the hierarchy itself will be illustrated by making specific reference to Figures 6 and 7. One interesting property is that faults propagate upward in the hierarchy. Thus, referring to Figure 7, if there is a leak in the vat, this causes a disturbance in the heat and water store, which in turn is reflected by changes in the mass and energy inventories, ultimately jeopardizing the achievement of the higher level goals. Reasons for proper system functioning are derived top-down in the hierarchy. To take two examples, the heater element exists to supply heat, and the mass and energy topology has been designed to satisfy the functional purposes of the system. Finally, the relations within levels of abstraction shown in Figure 6 provide information as to what will happen when certain changes are made (e.g., increasing the mass source flow rate will increase the mass inventory gradient).

To summarize, the means-end hierarchy representation of DURESS illustrates the various layers of functional structure that are inherent in the domain. Since it describes the way the system actually functions, this representation can be viewed as a normative model of the system (i.e., it specifies the way the operator should think of the system during problem solving), thereby providing a fundamental basis for interface design.

A. Mass Balance

1.) State Equation

$$\frac{dV}{dt} = \frac{E_0}{\rho} - \frac{V_0}{\rho}$$

2.) Geometrical

3.) Mapping

$$\frac{dV}{dt} = \frac{E_0}{\rho} - \frac{V_0}{\rho}$$

Below the equation, the variables are mapped to the geometric parameters:

- $V \rightarrow x$
- $V \rightarrow y$
- $\rho \rightarrow c$
- $E_0 \rightarrow a$
- $V_0 \rightarrow b$
- $\rho \rightarrow \theta$

Legend	
$V \rightarrow x$	$\rho \rightarrow c$
$V \rightarrow y$	$E_0 \rightarrow a$
$V \rightarrow \theta$	$V_0 \rightarrow b$
$\rho \rightarrow c$	
$E_0 \rightarrow a$	
$V_0 \rightarrow b$	
$\theta \rightarrow \rho$	

Figure 8. Making visible the invisible: mapping the relational structures of the work domain onto the geometrical properties of the interface.

Making the Affordances Visible

The next step in building an EID interface is to embed the means-end hierarchy structure into the interface. One can think of this phase as making visible the otherwise unobservable goal-relevant properties of the domain. Whereas traditional ecological optics is concerned with analyzing the information in the environment underlying the evolution of the visual system, the goal here is to reveal the process constraints in a form that is directly compatible with the visual system. According to the principles of EID, this is accomplished by mapping the relational structures of the domain, represented by the system's dynamic equations, onto the visible geometric properties of the interface objects. Figure 8 illustrates how this has been accomplished for DURESS.

In Figure 8A, the state equation describing the mass balance is given. This equation describes a relationship between a set of variables, and thus can be considered a form of domain knowledge. Following the goal of making visible the invisible, a geometric figure was developed to represent perceptually the relational structures specified by the state equation. As a result, there is an isomorphic mapping between the geometry of the trapezoid and the relationships between the variables describing the mass balance. Since the general form of the energy balance equation is very similar to that of the mass balance equation, a mapping similar to that shown in Figure 8A can also be developed for the energy balance.

Figure 8B represents the algebraic equation relating the temperature of the water in the reservoir, the energy contained therein, and the reservoir volume. Again, a geometric figure was developed to illustrate the relationships specified by the equation. The result is an isomorphic mapping between the trigonometric relationships describing the geometry of the triangle and the algebraic equation describing the behavior of the system.

The key behind these mappings is that they translate the constraints inherent in the domain into constraints on the interface geometry. This is equivalent to embedding knowledge about the domain in the interface. The great advantage of providing an externalized representation of the domain constraints in the interface is that it unburdens the operator's memory (Kotovsky, Hayes, and Simon, 1985). For example, rather than having to remember and retrieve the relationship describing the mass balance from memory, the operator can rely on the interface to provide this information. Making visible the invisible, then, allows operators to control the system by relying more on the powerful capabilities of perception and action than on comparatively more effortful and error-prone inferential processes.

The product of the design process up until this point is shown in Figure 9. Information at every level of the means-end hierarchy is included. Beginning at

the top level of Functional Purpose, the demand (D1, D2) and temperature (T1, T2) setpoints are represented in the interface. For the temperature settings, the upper and lower limits around the setpoints (40 °C and 20 °C) are shown as vertical lines on the two temperature scales (T1 and T2, respectively). The level of Abstract Function is represented by the group of graphics on the right. This portion of the display will be discussed in greater detail below. At the level of Generalized Function, the flowrates in each feedwater stream (e.g., FVA, FPA, FA1, FA2) and the heating rates (e.g., HTR1) are displayed as bar scales. At the level of Physical Function, the valve settings (e.g., VB) and heater settings (e.g., HTR2) are indicated by the small triangular pointers on the respective scales. Since the pump settings (e.g., PB) are discrete (either ON or OFF), they are directly labelled on the pumps themselves. Finally, at the level of Physical Form, the relative spatial layout of the components and the connections between them are also represented.

Figure 9. An EID interface for DURESS. See text for description.

With regard to the level of Abstract Function, the rectangular graphic on the left represents the mass balance for the reservoir, while the graphic on the right represents the energy balance. Both operate in a similar manner, as defined by

the mapping function in Figure 8A. Referring to Reservoir 1, the various inputs are shown at the top (e.g., MI1 for the mass and EI1 for the energy), the inventories on the side (e.g., V1 for volume, or mass, and E1 for energy), and the outputs at the bottom (e.g., D1 for demand, or mass, and EO1 for energy). The energy inputs (EI1 and EI2) are partialled out according to the two contributors. Thus, the energy added by the feedwater is shown as the lightly shaded bar, while the energy added by the heater is shown as the dark bar.

Intuitively, these energy and mass graphics rely on a funnel metaphor. Thus, if the bottom is wider than the top (i.e., output greater than input, as is the case with the mass balance for Reservoir 1 in Figure 9), then it is easy to visualize the consequence, namely that the volume should be decreasing. Thus, the slope of the line represents the rate at which the mass (or energy) inventory is changing (see Figure 8A). If input equals output, then the line would be perpendicular, indicating that the inventory should not change.

The graphic in the middle, between the mass and energy balances, illustrates the structure of the relationship between volume, energy, and temperature, following the mapping in Figure 8B. The horizontal line with a ball on the end that emanates from the current volume level is of fixed length. The height of this line always accompanies any change in volume (i.e., the bar will always be at the same height as the water level, V1 or V2). The thick diagonal line in the center display is always tangent to the ball on the edge of the horizontal line. Thus, a change in the vertical position of the horizontal line serves to change the slope of the line in the center display. For example, if volume increases, the horizontal line goes up, causing the diagonal to rotate counterclockwise, and thereby increasing the slope of the diagonal line. The slope of the diagonal represents the function that maps the amount of energy onto temperature (see Figure 8B). This mapping is indicated by the line from the energy inventory (E1, E2) that comes across and reflects off the diagonal and down onto temperature (T1, T2).

To summarize, the interface in Figure 9 has exploited the mappings described in Figure 8 to make all of the process dynamics visible to the operator. Whereas the interface in Figure 5 included information at the levels of Functional Purpose, Physical Function, and Physical Form, the EID interface in Figure 9 also includes the levels of Generalized Function and Abstract Function, thereby representing the entire means-end hierarchy.

Coupling Perception and Action

The final principle of EID is concerned with action. What means does the interface provide for the operator to control the system? EID suggests that, whenever possible, commands should be communicated by directly acting on

the display - the familiar idea of direct manipulation (Shneiderman, 1983). The intent here is to close the perception-action loop, thereby taking advantage of skilled motor control. As an example, to change the setting of the components in the EID interface in Figure 9 (e.g., VA, or HTR1), it is sufficient to "drag" the setpoints (represented as small triangles along the scales for the heaters and valves) to the new value using a mouse or a trackball. For the pumps, a change in state is achieved by merely clicking on the pump icon itself. Adopting a direct manipulation design is preferable to having the operator giving commands via abstract, and often awkward, command languages, as is common in some word processors.

In more complex systems, it may also be necessary to have a direct mapping between the perceptual hierarchy of nested display elements and the hierarchy of aggregated patterns of movement that is characteristic of skilled behavior (cf. Vicente and Rasmussen, 1989). An example of a display with these properties is discussed in Rasmussen and Vicente (in press).

A Few Scenarios

To give a more intuitive feel for how the interface would function dynamically, consider how several situations reveal themselves in the EID interface for DURESS. Take the simple case of a fault in one of the pumps. For example, what would happen if pump VB2 were to become blocked? This fault would be easy to diagnose with the interface in Figure 9, since there would be no flow through the valve (i.e., FB2 would go to zero) in spite of the fact that the valve setting, VB2, is open.

The higher order structure in the EID interface also allows operators to focus their attention by zooming in on the relevant part of the system. For instance, say there was an unexpected deviation in the temperature of Reservoir 1 (T1). If this deviation was due to a change in the slope of the diagonal line, then the operator should focus on the mass balance. If the diagonal did not change, then the operator could focus on the energy balance. Let us assume that the latter was the case. Now the question becomes: Was the unexpected change in energy inventory due to a change in the inputs (EI1) or the output (EO1)? Again, the structure of the information in the interface allows one to branch to the relevant subset of the system. Assuming that the change was in the inputs, one can then also determine which of the two inputs was responsible for the change.

The line of reasoning that was followed for this hypothetical scenario exploits the constraints represented by the means-end hierarchy in Figure 7. The fact that all of this information is represented in the interface allows operators to consider the system at a high level where the system description is simpler

(e.g., whether system goals are currently achieved or not), and then gradually navigate down the hierarchy to identify the cause of the abnormality. This example illustrates how the means-end hierarchy provides a mechanism for coping with the complexity of the domain.

Finally, the EID interface should also serve an instructional role by helping operators learn about the characteristics of the system. For instance, one of the properties that needs to be taken into account in controlling DURESS is the unidirectional coupling between the volume and temperature goals. As mentioned earlier, changing volume may result in an inadvertent change in temperature, but the reverse is not true. Thus, it would be best to concentrate on the mass balance first so as to stabilize the volume, and then worry about the temperature goal. If one were to start with the energy balance and get the temperature at the setpoint, and then go on to manipulate volume, the latter would move the temperature from the desired value. The operator would then have to go back to the heater to get the temperature back to the setpoint again. In an empirical study with a structurally similar system, Moray, Lootsteen, and Pajak (1986) found that most subjects caught on to the strategy of concentrating on volume first after 12 trials of controlling the system.

With the EID interface, operators should pick up on this more quickly. Stabilizing the mass balance first will fix the slope of the thick diagonal line in the center. Achieving the correct temperature is then easily accomplished by manipulating the heater until the energy inventory (E1, or E2) bounces off the diagonal and onto the acceptable region on the temperature scale. On the other hand, achieving the temperature setpoint first, and then going on to work on the mass balance will result in a change in volume, which in turn, would change the slope of the diagonal. This would then cause the temperature to deviate from the setpoint, requiring the operator to manipulate the heater once again to get the appropriate temperature. The fact that the domain constraints are revealed in the interface should make it easier for the operator to pick up on the structural system properties which need to be taken into account for effective control.

This discussion of how the EID interface for DURESS would react in different types of scenarios shows how this interface provides a much richer source of information than the interface illustrated in Figure 5.

CONCLUSION

In this paper, we have attempted to illustrate, first, how EID is consistent with the premises of ecological psychology, and second, how the EID framework can be put to practical use to develop a final design product. While there are many

reasons to believe that the ecological interface illustrated in Figure 9 represents a considerable improvement over traditional interfaces, only an experimental evaluation can confirm this intuition. An evaluation is currently being conducted to provide an empirical test of a subset of the EID framework within the context of DURESS (cf. Vicente, Flach, and Sanderson, 1989).

A question which remains to be addressed is the relevance and applicability of these ideas to real, complex human-machine systems. EID is directed at real-world problems, and therefore, it is critical that its principles generalize beyond domains that are used in laboratory research. In this regard, it is encouraging to note that Beltracchi (1987, 1989) has developed an overview display for nuclear power plants, based on the Rankine cycle heat engine, that is consistent with the ideas behind EID. The fact that this display has been implemented in an actual plant and has been positively received by operators (Lindsay and Staffon, 1988), suggests that the general principles of EID can be instantiated in real complex domains to produce effective mediators between operators and their work environment.

It is important to realize that EID is directed at a significant and pressing applied problem. While one would think that the problem of interface design for complex systems would have been "solved" by this point, it has actually yet to be addressed in an effective manner. Recently, the U. S. Nuclear Regulatory Commission convened a workshop, attended by a number of experts in several fields, to identify important research issues associated with interface design for complex human-machine systems (DeBor and Swezey, 1989). Several of the research topics that were identified as being of general and immediate importance are directly addressed by the research program centered around the EID framework (cf. Vicente, Flach, and Sanderson, 1989). Importantly, EID is a principled approach to the problem, since it has foundations in basic research in both cognitive engineering (Rasmussen, 1974, 1983, 1985a, 1986) and basic psychology (Gibson, 1966, 1979). The EID framework thus represents a coherent approach to a significant practical problem.

On a more general note, this paper has attempted to point to some of the similarities between cognitive engineering and ecological psychology. Interestingly, the parallels between the two disciplines can be traced to a common metatheoretical orientation. The ecological approach has many similarities with systems theory (cf. Lombardo, 1987, pp. 327-329). Gibson (1979, p. 2) himself explicitly referred to this connection. The discipline of cognitive engineering has been developed within a systems engineering framework as well (cf. Hollnagel and Woods, 1983; Norman, 1981; Rasmussen, 1986). Thus, even though the two disciplines are concerned with a very different set of problems, the way in which they approach those problems is

very similar, one could almost say identical. To the extent that the problems of the two areas share common generic solutions, the disciplines will benefit from interacting with one another.

The message that we would like to leave in closing is that the set of issues that is faced in designing information systems for complex work domains provides a mine of challenges, as well as opportunities, for ecological psychology (cf. Flach, 1989, in press b). There is one point in particular that is worthwhile emphasizing, since it has the greatest potential for influencing ecological psychology. There are certain classes of problems that must be continually faced in complex work domains which simply do not exist in navigating and locomoting in the natural environment. As mentioned in the introduction, the characteristics of these complex systems impose a set of demands that require that the operator engage in problem solving (i.e., going beyond the information given). However, as Jenkins (1989) has pointed out, the ecological community has tended to focus on perception and action to the detriment of cognition. A greater emphasis on the study of cognition seems warranted if the ecological approach is to provide a comprehensive account of human behavior.

The key point is that the type of applied problems discussed in this paper can encourage such a development. Given the overlap in perspectives that already exists between ecological psychology and cognitive engineering, complex work domains offer themselves as a natural, yet challenging, testbed for posing basic research questions, which can, in turn, lead to new developments in ecological psychology. Therefore, just as Gibson's involvement with the applied concerns of driving (Gibson and Crooks, 1938) and flying (Gibson, 1947) led him to a revolutionary theory of perception (cf. Lombardo, 1987; Reed, 1988), the challenges of complex work domains can also serve as the catalyst for a new ecological theory of cognition.

There is little doubt that an interaction between the ecological psychology and cognitive engineering communities is bound to be fruitful. While there are a few small steps being taken in this direction (e.g., Flach, 1989, in press b; Flach and Vicente, 1989; Kirlik, 1989; Kugler, 1989; Kugler and Lintern, 1989; Moray, 1989; Rasmussen and Vicente, 1989; Vicente and Rasmussen, 1989; Woods, 1989), for the most part, cognitive engineering is still an unexplored ecological frontier.

ACKNOWLEDGEMENTS

This paper is based on a paper (Rasmussen and Vicente, in press) and a poster that were presented at the symposium on The Ecology of Human-Machine

Systems during the Fifth International Conference on Event Perception and Action held at Miami University, Oxford, Ohio on July 26, 1989.

Many thanks are due to: Jan Kofoed for developing the DURESS simulation; Suzanne Faikus for doing an excellent job of designing the figures; and John Flach, James Jenkins, Alex Kirlik, Peter Kugler, Bill Mace, Gary Riccio, Sverker Runeson, David Woods, Leo Beltracchi, Len Goodstein, Penny Sanderson, Bob Shaw, and Michael Turvey for stimulating discussions and insightful comments.

REFERENCES

- Abraham, R. H., and Shaw, C. D. (1982). Dynamics - The geometry of behavior, part 1: periodic behavior. Santa Cruz, CA: Aerial Press.
- Abraham, R. H., and Shaw, C. D. (1983). Dynamics - The geometry of behavior, part 2: chaotic behavior. Santa Cruz, CA: Aerial Press.
- Abraham, R. H., and Shaw, C. D. (1984). Dynamics - The geometry of behavior, part 3: global behavior. Santa Cruz, CA: Aerial Press.
- Abraham, R. H., and Shaw, C. D. (1988). Dynamics - The geometry of behavior, part 4: bifurcation behavior. Santa Cruz, CA: Aerial Press.
- Allen, T. F. H., and Starr, T. B. (1982). Hierarchy: Perspectives for ecological complexity. Chicago: University of Chicago Press.
- Ashby, W. R. (1956). An introduction to cybernetics. London: Methuen.
- Beltracchi, L. (1987). A direct manipulation interface for water-based rankine cycle heat engines. IEEE Transactions on Systems, Man, and Cybernetics, SMC-17, 478-487.
- Beltracchi, L. (1989). Energy, mass, model-based displays, and memory recall. IEEE Transactions on Nuclear Science, 36, 1367-1382.
- Bernstein, N. (1967). The coordination and regulation of movements. Oxford: Pergamon Press.
- Blattner, M. A., Sumikawa, D. A., and Greenberg, R. M. (1989). Earcons and icons: Their structure and common design principles. Human-Computer Interaction, 4, 11-44.
- Bødker, S. (1989). A human activity approach to user interfaces. Human-Computer Interaction, 4, 171-195.
- Crossman, E. R. F. W., and Cooke, J. E. (1974). Manual control of slow-response systems. In E. Edwards and F. P. Lees (Eds.), The human operator in process control (pp. 51-66). London: Taylor and Francis.
- DeBor, J., and Swezey, R. (1989). Man-machine interface issues in nuclear power plants (NUREG/CR-5348). Washington, DC: USNRC.
- Flach, J. M. (1989). An ecological alternative to egg-sucking. Human Factors Society Bulletin, 32(9), 4-6.
- Flach, J. M. (in press a). Control with an eye for perception: Precursors to an active psychophysics. Ecological Psychology.
- Flach, J. M. (in press b). The ecology of human-machine systems I: Introduction. Ecological Psychology.
- Flach, J. M., and Vicente, K. J. (1989). Complexity, difficulty, direct manipulation and direct perception. Manuscript submitted for publication.
- Gaver, W. W. (1986). Auditory icons: Using sound in computer interfaces. Human-Computer Interaction, 2, 167-177.
- Gaver, W. W. (1989). The SonicFinder: An interface that uses auditory icons. Human-Computer Interaction, 4, 67-94.
- Gibson, J. J. (1947). Motion picture testing and research. Aviation Psychology Research Reports, No. 7. Washington: U.S. Government Printing Office.

- Gibson, J. J. (1961). Ecological optics. Vision Research, 1, 253-262.
- Gibson, J. J. (1966). The senses considered as perceptual systems. Boston: Houghton- Mifflin.
- Gibson, J. J. (1979). The ecological approach to visual perception. Boston: Houghton- Mifflin.
- Gibson, J. J. (1982a). Notes on affordances. In E. Reed and R. Jones (Eds.), Reasons for realism: Selected essays of James J. Gibson (pp. 401-418). Hillsdale, NJ: LEA.
- Gibson, J. J. (1982b). Notes on direct perception and indirect apprehension. In E. Reed and R. Jones (Eds.), Reasons for realism: Selected essays of James J. Gibson (pp. 289-293). Hillsdale, NJ: LEA.
- Gibson, J. J., and Crooks, L. E. (1938). A theoretical field analysis of automobile-driving. American Journal of Psychology, 51, 453-471.
- Goodstein, L. P. (1981). Discriminative display support for process operators. In J. Rasmussen and W. B. Rouse (Eds.), Human detection and diagnosis of system failures (pp. 433-449). New York: Plenum.
- Goodstein, L. P. (1983). An integrated display set for process operators. In G. Johannsen and J. E. Rijnssdorp (Eds.), Analysis, design and evaluation of man-machine systems (pp. 63-70). Oxford: Pergamon.
- Goodstein, L. P., Andersen, H. B., and Olsen, S. E. (1988). Tasks, errors, and mental models: A festschrift to celebrate the 60th birthday of Professor Jens Rasmussen. London: Taylor and Francis.
- Greene, P. H. (1982). Why is it easy to control your arms? Journal of Motor Behavior, 14, 260-286.
- Hollnagel, E. (1983). What we do not know about man-machine systems. International Journal of Man-Machine Studies, 18, 135-143.
- Hollnagel, E., and Woods, D. D. (1983). Cognitive systems engineering: New wine in new bottles. International Journal of Man-Machine Studies, 18, 583-600.
- Hollnagel, E., Mancini, G., and Woods, D. D. (1986). Intelligent decision support in process environments. Berlin: Springer-Verlag.
- Hollnagel, E., Mancini, G., and Woods, D. D. (1988). Cognitive engineering in complex dynamic worlds. London: Academic Press.
- Jenkins, J. (1989). An ecological orientation to cognition. Paper presented at the 5th International Conference on Event Perception and Action. Oxford, OH.
- Kirlik, A. C. (1989). The organization of perception and action in complex control skills (Unpublished Doctoral Dissertation). Columbus, OH: The Ohio State University, Department of Industrial and Systems Engineering.
- Korf, R. E. (1987). Planning as search: A quantitative approach. Artificial Intelligence, 33, 65-88.
- Kotovsky, K., Hayes, J. R., and Simon, H. A. (1985). Why are some problems hard? Evidence from tower of Hanoi. Cognitive Psychology, 17, 248-294.
- Kugler, P. N. (1989). Of course it can: The ecology of workspaces. Paper presented at the 5th International Conference on Event Perception and Action. Oxford, OH.
- Kugler, P. N., and Lintern, G. (1989). Self-organization and the evolution of instability: The awakening of sleeping nonlinearities. In Proceedings of the World Bank Conference on Safety and Risk Management. Washington, DC: World Bank.
- Kugler, P. N., and Turvey, M. T. (1987). Information, natural law, and the self-assembly of rhythmic movement. Hillsdale, NJ: LEA.
- Laurel, B. K. (1986). Interface as mimesis. In D. A. Norman and S. W. Draper (Eds.), User centered system design: New perspectives on human-computer interaction (pp. 67-85). Hillsdale, NJ: LEA.
- Lindsay, R. W., and Staffon, J. D. (1988). A model based display system for the experimental breeder reactor-II. Paper presented at the Joint Meeting of the American Nuclear Society and the European Nuclear Society. Washington, DC.
- Lombardo, T. J. (1987). The reciprocity of perceiver and environment: The evolution of James J. Gibson's ecological psychology. Hillsdale, NJ: LEA.

- Mace, W. M. (1980). Perceptual activity and direct perception. Behavioral and Brain Sciences, 3, 392-393.
- Meister, D. (1985). Behavioral analysis and measurement methods. New York: Wiley.
- Mesarovic, M. D., Macko, D., and Takahara, Y. (1970). Theory of hierarchical, multilevel, systems. New York: Academic Press.
- Michaels, C. F., and Carello, C. (1981). Direct perception. New York: Prentice-Hall.
- Mitchell, C. M., and Miller, R. A. (1986). A discrete control model of operator function: A methodology for information display design. IEEE Transactions on Systems, Man, and Cybernetics, SMC-16, 343-357.
- Mitchell, C. M., and Saisi, D. L. (1987). Use of model-based qualitative icons and adaptive windows in workstations for supervisory control systems. IEEE Transactions on Systems, Man, and Cybernetics, SMC-17, 573-593.
- Moray, N. (1989). Designing human-machine systems: Can ecological psychology help? Paper presented at the 5th International Conference on Event Perception and Action. Oxford, OH.
- Moray, N., Lootsteen, P., and Pajak, J. (1986). Acquisition of process control skills. IEEE Transactions on Systems, Man, and Cybernetics, SMC-16, 497-504.
- Newman, J. R. (1966). Extension of human capability through information processing and display systems (SP-2560). Santa Monica, CA: System Development Corporation.
- Norman, D. A. (1981). Steps toward a cognitive engineering: System images, system friendliness, mental models. Paper presented at Symposium on Models of Human Performance, ONR Contractor's Meeting. La Jolla, CA: University of California at San Diego.
- 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. 32-61). Hillsdale, NJ: LEA.
- Pattee, H. H. (1972). The nature of hierarchical controls in living matter. In R. Rosen (Ed.), Foundations of mathematical biology, vol. 1 - Subcellular systems (pp. 1-21). New York: Academic Press.
- Pattee, H. (1986). Universal principles of measurement and language functions in evolving systems. In J. L. Casti and A. Karlqvist (Eds.), Complexity, language, and life: Mathematical approaches (pp. 268-281). Berlin: Springer-Verlag.
- Perrow, C. (1984). Normal accidents: Living with high risk technologies. New York: Basic Books.
- Rasmussen, J. (1974). The human data processor as a system component: Bits and pieces of a model (Risø-M-1722). Roskilde, Denmark: Danish Atomic Energy Commission.
- Rasmussen, J. (1983). Skills, rules, and knowledge; Signals, signs, and symbols, and other distinctions in human performance models. IEEE Transactions on Systems, Man, and Cybernetics, SMC-13, 257-266.
- Rasmussen, J. (1985a). The role of hierarchical knowledge representation in decisionmaking and system management. IEEE Transactions on Systems, Man, and Cybernetics, SMC-15, 234-243.
- Rasmussen, J. (1985b). Conceptual models in man-machine design verification (Risø-M-2520). Roskilde, Denmark: Risø National Laboratory, Electronics Department.
- Rasmussen, J. (1986). Information processing and human-machine interaction: An approach to cognitive engineering. New York: North-Holland.
- Rasmussen, J. (1988a). Information technology: A challenge to the Human Factors Society? Human Factors Society Bulletin, 31(7), 1-3.
- Rasmussen, J. (1988b). A cognitive engineering approach to the modeling of decision making and its organization in: Process control, emergency management, cad/cam, office systems, and library systems. In W. B. Rouse (Ed.), Advances in man-machine systems research, vol. 4 (pp. 165-243). Greenwich, CT: JAI Press.
- Rasmussen, J., and Goodstein, L. P. (1988). Information technology and work. In M. Helander (Ed.), Handbook of human-computer interaction (pp. 175-201). New York: North Holland.
- Rasmussen, J., and Rouse, W. B. (1981). Human detection and diagnosis of system failures. New York: Plenum.

- Rasmussen, J., and Vicente, K. J. (1989). Coping with human errors through system design: Implications for ecological interface design. International Journal of Man-Machine Studies, 31, 517-534.
- Rasmussen, J., and Vicente, K. J. (in press). Cognitive engineering: An ecological frontier (Risø Tech. Rep.). Roskilde, Denmark: Risø National Laboratory, Department of Information Technology.
- Rasmussen, J., Duncan, K., and Leplat, J. (1987). New technology and human error. New York: Wiley.
- Reed, E. S. (1988). James J. Gibson and the psychology of perception. New Haven, CT: Yale University Press.
- Reed, E., and Jones, R. (1982). Reasons for realism: Selected essays of James J. Gibson. Hillsdale, NJ: LEA.
- Rosen, R. (1978). Fundamentals of measurement and representation of natural systems. New York: North-Holland.
- Runeson, S. (1977). On the possibility of "smart" perceptual mechanisms. Scandinavian Journal of Psychology, 18, 172-179.
- Saltzman, E. L., and Munhall, K. G. (1989). A dynamical approach to gestural patterning in speech production. Ecological Psychology, 1, 333-382.
- Shaw, R., and Bransford, J. (1977). Introduction: Psychological approaches to the problem of knowledge. In R. Shaw and J. Bransford (Eds.), Perceiving, acting and knowing: Toward an ecological psychology (pp. 1-39). Hillsdale, NJ: LEA.
- Shaw, R. E., and Turvey, M. T. (1981). Coalitions as models for ecosystems: A realist perspective on perceptual organization. In M. Kubovy and J. R. Pomerantz (Eds.), Perceptual organization (pp. 343-415). Hillsdale, NJ: LEA.
- Shaw, R., Turvey, M. T., and Mace, W. (1982). Ecological psychology: The consequence of a commitment to realism. In W. B. Weimer and D. S. Palermo (Eds.), Cognition and the symbolic processes, vol. 2 (pp. 159-26). Hillsdale, NJ: LEA.
- Sheridan, T. B. (1987). Supervisory control. In G. Salvendy (Ed.), Handbook of human factors (pp. 1243-1268). New York: Wiley.
- Sheridan, T. B., and Johannsen, G. (1976). Monitoring behaviour and supervisory control. New York: Plenum.
- Shneiderman, B. (1983). Direct manipulation: A step beyond programming languages. IEEE Computer, 16(8), 57-69.
- Simon, H. A. (1981). The sciences of the artificial (2nd ed.). Cambridge, MA: MIT Press.
- Turvey, M. T. (1977). Preliminaries to a theory of action with reference to vision. In R. Shaw and J. Bransford (Eds.), Perceiving, acting and knowing: Toward an ecological psychology (pp. 211-265). Hillsdale, NJ: LEA.
- Turvey, M. T., and Kugler, P. N. (1984). A comment on equating information with symbol strings. American Journal of Physiology, 246, R925-927.
- Turvey, M. T., and Shaw, R. (1979). The primacy of perceiving: An ecological reformulation of perception for understanding memory. In L.-G. Nilsson (Ed.), Perspectives on memory research: Essays in honor of Uppsala University's 500th anniversary (pp. 167-222). Hillsdale, NJ: LEA.
- Turvey, M. T., Shaw, R. E., and Mace, W. (1978). Issues in the theory of action: Degrees of freedom, coordinative structures and coalitions. In J. Requin (Ed.), Attention and performance VII (pp. 557-595). Hillsdale, NJ: LEA.
- Turvey, M. T., Shaw, R. E., Reed, E. S., and Mace, W. M. (1981). Ecological laws of perceiving and acting: In reply to Fodor and Pylyshyn (1981). Cognition, 9, 237-304.
- Veldhuyzen, W., and Stassen, H. G. (1976). The internal model: What does it mean in human control? In T. B. Sheridan and G. Johannsen (Eds.), Monitoring behaviour and supervisory control (pp. 157-170). New York: Plenum.
- Vicente, K. J. (1987). DURESS: A domain for cognitive engineering research. Unpublished manuscript.

- Vicente, K. J. (1988). Adapting the memory recall paradigm to evaluate interfaces. Acta Psychologica, 69, 249-278.
- Vicente, K. J. (1990). Ecological interface design as an analytical evaluation tool. In ANS Topical Meeting on Advances in Human Factors Research: Nuclear and Beyond. Nashville, TN: American Nuclear Society.
- Vicente, K. J., and Rasmussen, J. (1989). Ecological interface design: Theoretical foundations. Manuscript submitted for publication.
- Vicente, K. J., Flach, J. M., and Sanderson, P. M. (1989). Ecological interface design: Creating a window to the world of complex work domains. In Proceedings of the 22nd Annual Meeting of the Human Factors Association of Canada (pp. 151-155). Toronto: HFAC/ACE.
- Warren, W. H., Jr. (1985). Environmental design as the design of affordances. Paper presented at the 3rd International Conference on Event Perception and Action. Uppsala, Sweden.
- Wickens, C. D. (1986). The effects of control dynamics on performance. In K. R. Boff, L. Kaufman, and J. P. Thomas (Eds.), Handbook of perception and human performance, vol. II (Ch. 39). New York: Wiley.
- Woods, D. D. (1984). Visual momentum: A concept to improve the coupling of person and computer. International Journal of Man-Machine Studies, 21, 229-244.
- Woods, D. D. (1986). Paradigms for intelligent decision support. In E. Hollnagel, G. Mancini, and D. D. Woods (Eds.), Intelligent decision support in process environments (pp. 153-173). Berlin: Springer-Verlag.
- Woods, D. D. (1988). Coping with complexity: The psychology of human behaviour in complex systems. 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. 128-148). London: Taylor and Francis.
- Woods, D. D. (1989). The ecology of problem solving. Paper presented at the 5th International Conference on Event Perception and Action. Oxford, OH.
- Woods, D. D. (in press). The cognitive engineering of problem representations. In J. Alty and G. Weir (Eds.), Human-computer interaction and complex systems. London: Academic Press.
- Woods, D. D., and Elias, G. (1988). Significance messages: An integral display concept. In Proceedings of the Human Factors Society 32nd Annual Meeting (pp. 1350-1354). Santa Monica, CA: HFS.
- Woods, D. D., and Hollnagel, E. (1987). Mapping cognitive demands in complex problem-solving worlds. International Journal of Man-Machine Studies, 26, 257-275.
- Woods, D. D., and Roth, E. M. (1988). Cognitive systems engineering. In M. Helander (Ed.), Handbook of human-computer interaction (pp. 3-41). New York: North Holland.
- Woods, D. D., O'Brien, J. F., and Hanes, L. F. (1987). Human factors challenges in process control: The case of nuclear power plants. In G. Salvendy (Ed.), Handbook of human factors (pp. 1724-1770). New York: Wiley.