5 Ecological Interface Design

Introduction

The practice of ecological interface design (EID) originated in the work of Jens Rasmussen, who spent much of his career as a cognitive systems engineer for the Halden Nuclear Reactor Project in Denmark. Over the course of his work, he studied a range of issues, from the ways in which people diagnose faults in electrical circuits to how they monitor and control complex industrial processes (figure 5.1). Across this work he developed a theory of human activity (focusing primarily on industrial processes but extensible to any situation in which people interact with technology). As we shall see, this theory owes much to the concepts of Gibson; even if Rasmussen's theory does not directly draw on embodied cognition, there is much overlap and scope to build on the parallels.

Given the prevalence of digital technology, the design of the user interface (i.e., the manner in which information is presented to the person and the manner in which the person's activity is invited and supported by a device) is of central importance. In human-computer interaction (HCI), an influential account of the reciprocal relationship between technology and human activity is the task-artifact cycle[1] (figure 5.2).

The task-artifact cycle points to two core themes for HCI and for design more generally. The first is that the tasks that people perform are constrained by the device. The usefulness of this apparent truism becomes obvious when you realize that "task" is not simply the performance of an action but also the setting of a goal and definition of fitness (i.e., how well the goal is met). In this respect, the task-artifact cycle illustrates the ongoing reciprocal engagement in the human-artifact-environment system. It does so in terms

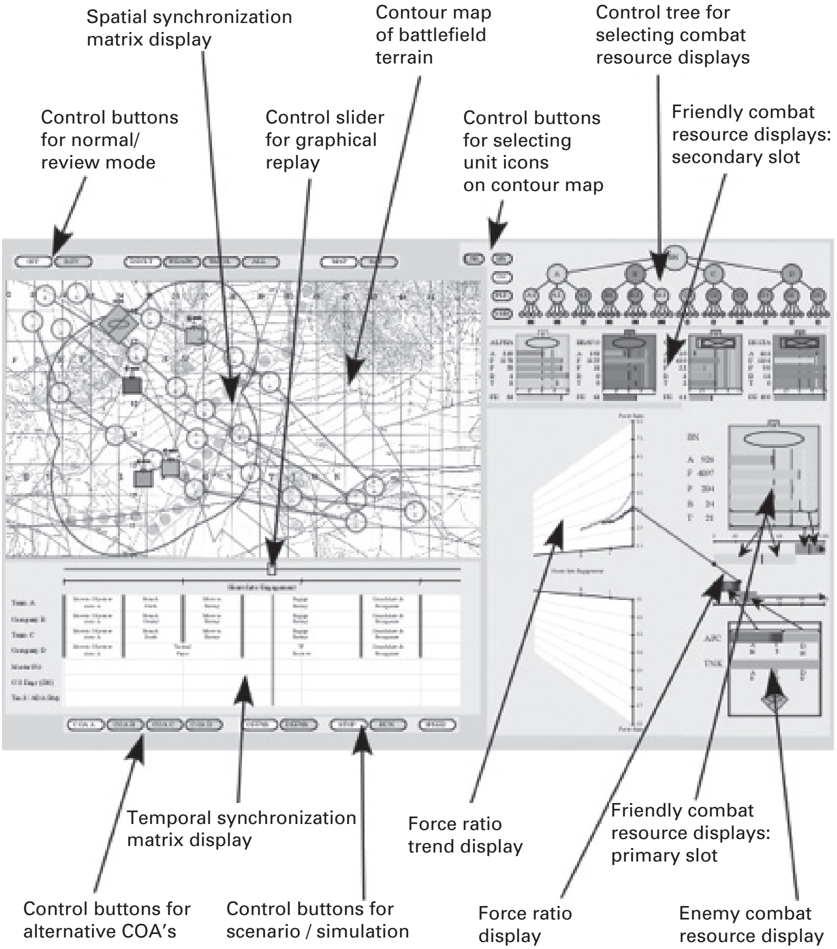

Figure 5.1
Example of an ecological interface.

of the relationship between human and artifact defined as action (task) and outcome (goal). While the environment is not explicitly included, one can expect the outcome to be further constrained by the environment. Further, the focus on "task" might imply solely physical action but it carries a richer meaning that includes the interpretation of the goal. To illustrate this, consider writing a letter with a typewriter (if you can remember such a device), or a word processing program on a laptop, or with a stylus on a tablet. There are obvious differences between these artifacts and in the physical

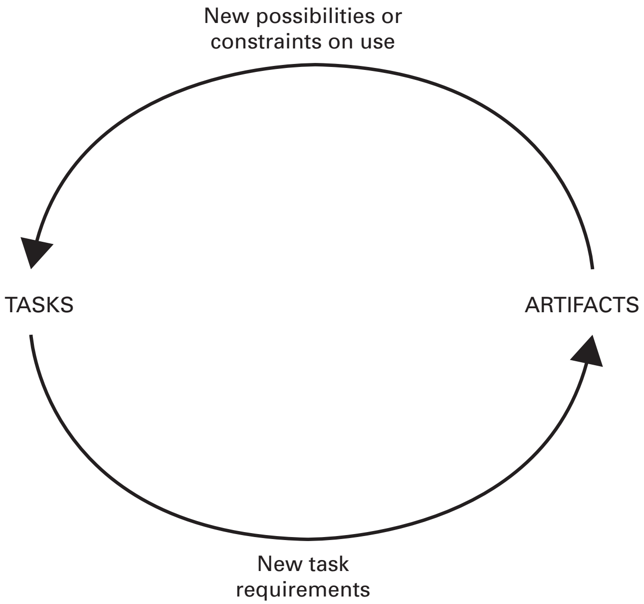

Figure 5.2
Task-artifact cycle.

actions that are made to form words and to manage the layout of the pages of text. In addition, there are differences in how mistakes might be corrected, which could have an impact on how much care you might apply in forming the words (particularly if, like the typewriter, the technology is unforgiving of error). In this case, the performance of each task could be made with greater or lesser awareness of contingent tasks (such as error correction). Moreover, the "outcome" of the activity (in terms of the content, layout, and appearance of the letter when it is printed) has a bearing on the performance of the tasks, both in terms of the meta-tasks associated with tabs, margins, spacing, and other settings for the page setup, the choice of font, and son on and the impact of these meta-tasks on the activity (e.g., should these meta-tasks be performed prior to typing the letter or can they be amended at any point during the typing?). These considerations point to the second theme, which is that (re)design of the artifacts can be based on problems that arise from these contingent and meta-tasks. As these "problems" become more clearly defined, so their "solution" leads to new versions of the artifact.

Depending on how one views this relationship between problems and solutions (in the development of technology), one can conceptualize the design of new artifacts either as a linear progression towards the "best" version or as a patchwork of changes. Although there is a sense in which the problem-solution pattern might imply a dialectical progression, few people subscribe to the idea that technology develops in a linear manner. Rather, most accounts of the evolution of technology recognize that developments are in response to specific problems (and different people using the same artifact for different tasks might encounter different problems). In an interesting account of technology development, Bijker[2] suggests that the recognition of "problems" with technology relate not only to the people using it (and their uses of it) but also to the "technological frame" that they apply to it. For designers, the "technological frame" relates to the tradition in which you have been working and to the manner in which this tradition can be applied, through analogy, to new problems. For example, in the development of the Internet, some researchers came from a telephone network tradition and emphasized the ways in which data could be combined into packets to route through the network; others came from the Air Defence Systems and emphasized the need for the network to be secure and robust; others came from traditions that emphasized human interactions with information.[3] Each tradition brought its own definition of concepts such as "network" or "information" and different expectations of how people would interact with these. One of the key aspects of the development of the Internet (in addition to the technical achievements) was agreement on which aspects of which definitions to include in order to reach consensus on what to build and how it should operate.

The task-artifact cycle illustrates the manner in which user activity is constrained by a given technology. This concept of constraint is not only critical to the definition of embodied cognition presented in this book, but also key to the approach of EID. In order to explain this approach, we need to appreciate the relationship between the idea of task ecology and cognitive work analysis.

Cognitive Work Analysis

The notion of task ecology was presented in chapter 3. Rasmussen proposed several ways to support a description of the ecology in which a system operates. For example, in AcciMaps[4] the system is decomposed into levels that range from the political/regulatory to specific activities performed by people. The aim of this description is to provide a framework within which to understand the interacting constraints on a system during its operation prior to an accident (hence, the “acci” in the title).

From the notion of task ecology developed in chapter 3, we can say that constraints, on the human-artifact-environment system can be defined at different levels, such that there is the level at which certain activities are impossible, there is the level at which certain activities prevent performance of other activities, there is the level at which certain activities could result in dangerous or undesirable outcomes, and there is the level at which some outcomes are valued more highly than others. In other words, constraints will be defined at different levels, from physical to social and economic, to political or societal. Rasmussen describes these in terms of whether responsibility for outcome lies with the operators, their manager, or the government agencies regulating the industry. In this broadening of the definition of a system, Rasmussen was considering the social implications of technology, with much the same motivation as socio-technical systems[5] or actor-network theory.[6] In Rasmussen’s approach, each level has a bearing on the salience of information (i.e., sets of features from the environment) or the choice of action or the definition of outcomes (in terms of acceptability or desirability). The different levels give rise to the idea of a part-whole configuration, which reflects the focus of attention (from physical to social) of the people acting in that environment. An obvious consequence of this idea, and one that features in accident investigations, is that undue focus on one level could result in failure to recognize, or misinterpret, constraints at another level. While I have used the word “levels” in my discussion of Rasmussen, it is important to note that he was more concerned with “transitions and overlaps” within a hierarchy of “units” (or views) of a work domain. Other methods developed by Rasmussen and his colleagues focused on the boundaries of the system and are couched in terms of what the system seeks to achieve and avoid—in other words, on the relationship between the functional purpose and desirable outcomes. This is presented in the form of a work domain analysis.

In the work domain analysis (figure 5.3) means-ends relations (as considered in terms of problem-solving in chapter 2) are captured in a hierarchy of five levels (from bottom to top): (1) the physical form, which

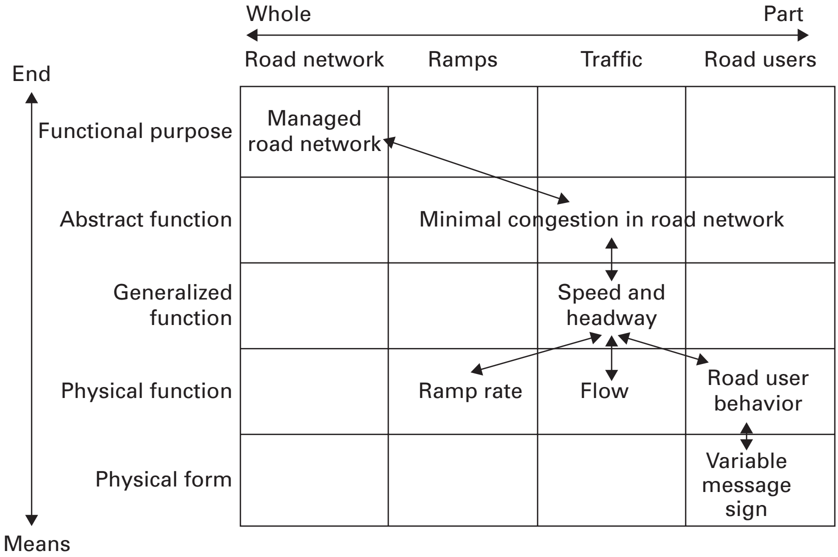

Figure 5.3
Work domain analysis.[7]

can be considered in terms of the features of the environment (in terms of the input to the system, e.g., control devices, information sources, and the like); (2) the physical functions that can performed by the system on the environment; (3) the generalized function, which can be considered in terms of the high-level goals which can satisfy the abstract functions; (4) the abstract function(s), which can be considered in terms of the desired outcomes ("values and priorities") that constrain system activity; and (5) the functional purpose which is the raison d'etre of the system.

It is not my intention here to provide guidance on how to create this diagram (but there is excellent guidance available for the methods that underpin this general approach of cognitive work analysis [CWA][8]). The "abstraction" reflects Rasmussen's observation that people (particularly those he studied who were fault-finding in electronic circuits or controlling power stations) engage in different forms of cognitive activity at different levels, from the immediate local features of the environment to more global concerns as to the functional purpose of the system. The "decomposition" reflects Rasmussen's observation that the different "abstraction" levels will define salience in different ways. This moving between levels of detail is also characteristic of domains such as intelligence analysis, where analysts appear to move between "broad" and "narrow" views of a situation.[9] Taken together, the notions of abstraction and decomposition suggest that problem-solving involves consideration of a system at different levels of purpose and in different levels of detail.

The abstraction-decomposition space can be "read" top-down (in terms of how the desired outcomes enable and constrain activities) and bottom-up (in terms of how the objects in the system support the performance of activities). From figure 5.3, the description is concerned with appreciating the "why" of the system (in terms of what it is seeking to achieve and how success is defined) and the "how" of the system (in terms of what activities contribute to success). In this way, the abstraction-decomposition space illustrates the relations between the physical function and functional purpose levels, and it *appears* to reflect the knowledge that is used in problem-solving. Having said this, I would add that it is not apparent that the abstraction-decomposition space or the analysis that produces it is capable of capturing "tacit" knowledge, so much as assuming that the analyst has produced a "formal" description, which returns us to the discussion of information-as-content versus information-as-context in the earlier chapters. This point will be explored further in the next section. For this section, a further point to note is that the abstraction-decomposition space implies that means-ends analysis relates to the connections between each level. This means that the definition of a "means" slips from an activity to a "function" in ways that are not obviously useful or ontologically reliable (similar problems relate to the conflation of "end" with "purpose"). This conflation of "means" with "function" can create problems, particularly for the novice user of the method and, as far as I am aware, is not an issue that the CWA community has fully resolved.

Before considering decision ladders, it is worth digressing slightly on the topic of "tasks." Ergonomics has long been concerned with defining the "tasks" that people perform in the workplace. In these terms, a "task" is a discrete action that is performed in order to achieve a goal. From the perspective of CWA, however, this notion is too restrictive because it implies that the task-goal mapping points to a "one-best-way" to do work (i.e., the criticism points to a Taylorist tendency to routinize work to the most efficient strategies). I do not agree with this interpretation (task analysis methods are quite capable of reflecting alternative strategies), and a "task" is essentially a system goal that needs to be achieved (this echoes the point made earlier regarding "tasks" in the task-artifact cycle).

Decision Ladders and Decision Strategies

We began this chapter with a consideration of "tasks" (in the task-artifact cycle) and the observation that, by implication, the "task" includes its goal and definition of fitness. Navigating the abstraction-decomposition space involves focusing on a task as a possible route within the space of possibilities—in other words, choosing a course of activity (defined as a "strategy" by which information is selected and actions are performed) that relates to a functional purpose. A strength of CWA (in comparison with most other methods) is that it can comfortably accommodate multiple alternative routes. This means that, rather than being bound to a specific instance of the system in a specific instance of an environment, the approach seeks to be event-independent. To do this, CWA uses a form of visualization known as the "decision ladder."

Schematically, the decision ladder (figure 5.4) represents activity as a simple, linear input-output flow. The input is a stimulus from information sources (described as physical objects in the abstraction hierarchy), and the output is an action. The decision ladder assumes that there is a canonical sequence of steps, which can be considered in terms of information-processing psychology. Diagrammatically, the flow is split into two "legs"; in the left-hand leg, the input leads up to the goal (purpose) as a process that identifies a discrepancy between input and desired state, and, in the right-hand leg, an action is specified that returns the system to its goal state. Each of the intervening steps is labeled in terms of the information sources that are required to complete the task. In essence, the process is one in which the environment is sampled and a step completed (either as an action to collect information or to process the information to pass to the next step). From figure 5.4, one clearly sees the means-ends approach to problem-solving and the information-processing tradition. The reader might be a little nonplussed as to why I am presenting this. The step ladder clearly shows the serial, "production line" of information-processing that I complained about in chapter 1. Also, it seems to reflect the standard operating procedure (SOP) by which activity is performed. As is well known, experience leads us to adopt all manner of paths to circumvent the SOP,

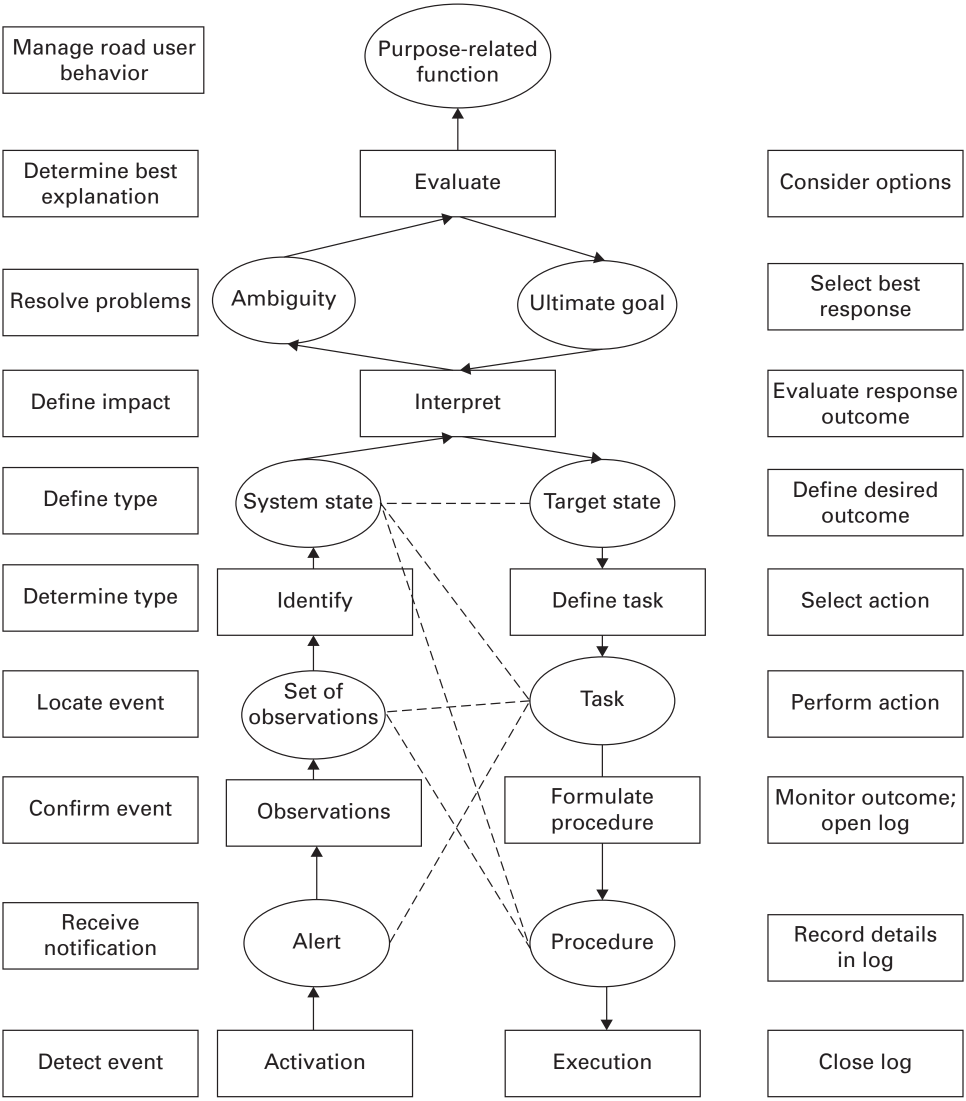

Figure 5.4
Decision ladder.

and it is these paths (shown as dotted lines on figure 5.4) that are intended to capture the strategies that are used to respond to different environmental situations or system states.

Lintern[10] proposes that the paths between each side of the step ladder can take three forms: state transitions (a term he prefers rather than "short-cuts"), which retain commitment to the SOP in that these might be performed under certain states; shunts, which involve deliberative transformation of one "state" to another; and leaps, in which there is a direct association between "states" with no intervening process. While I am using the word "states" here to refer to the human-artifact-environment system, it is not clear how Rasmussen uses the word. For Rasmussen, the decision ladder shows how "rational, causal reasoning connects the 'states of knowledge' in the basic sequence."[11] It is a moot point as to whether these "states of knowledge" (including the sequence itself) constitute a mental model (internal representation) that the person creates and maintains during their activity. An information-processing interpretation would assume this to be the case. Alternatively, Rasmussen saw the decision ladder as a "a map useful to represent the structure of such a model"[12] (which could, I guess, be read as the SOP by which a trained person might conduct their work). Yet, as we shall see later in this chapter, his use of ecological psychology could mean that the "states of knowledge" are just as likely to be realized with objects in the environment (or, at least, objects on a visual display). This seems to echo the notion of off-loading from distributed cognition and the separation of knowledge in the world and knowledge in the head.[13]

In terms of the strategies that can be used to navigate the space of possibilities, Rasmussen contrasted "skill-based," "rule-based," and "knowledge-based" action. In "Knowledge-based" action the problem solver seeks to define the problem and deduce the most appropriate sequence of actions by which to solve the problem. This follows from the information-processing approaches to problem-solving that we saw in chapter 2 and is intended to reflect the strategy that might be taken to deal with a novel problem. "Rule-based" action involves the application of a known strategy to solve the problem (either because the problem itself is familiar or because the problem is sufficiently analogous a familiar problem to allow rules to be transferred). Again, this approach was considered in chapter 2. Both of these approaches fall squarely within the remit of information-processing psychology. "Skill-based" action, on the other hand, owes a debt to Rasmussen's readings of Gibson's perception-action coupling. The second clause of the last sentence is deliberately clumsy because the relation between "skill-based" action and Gibsonian concepts is not as clear cut as some writers have claimed. However, the key point here is that the decision ladder represents a marriage between information-processing and "embodied" theories in ways that I have been trying to prevent.

Before exploring what an "embodied" reading of Rasmussen's thoughts might look like (that is, how one might read the word "ecological" in terms of radical embodied cognitive science [RECS]), it is worth noting comparisons that have been drawn between the decision ladder and recognition-primed decision-making (RPD) as descriptions of "expert" performance (see chapter 3). By decomposing activity into task steps, the decision ladder is concerned with defining what tasks need to be done while implicitly indicating that these tasks will be performed by one or other entity in the system, while RPD does not draw so clear a distinction between tasks (nor does it imply that the "system" involves more than one entity).[14] Further, RPD (in its assumption that the "expert" experiences the situation in terms of schema drawn from prior experience) proposes that there are features in the environment to which the "expert" can respond, while the decision ladder proposes that the salience of the features will be derived by the "expert" during their analytical activity.[15] For our purposes, RPD *could* be framed in "embodied" terms (chapter 3) rather than relying on the information-processing foundations of a decision ladder.

The decision ladder could be related to the task-artifact cycle in terms of ways in which to discover new possibilities or constraints on the information available to the system (which is also a core objective of CWA more broadly). In contrast, RPD (particularly in the "embodied" form we are advocating) is concerned with appreciating the ways in which particular sets of features in the environment acquire salience for the expert in order to better understand how to abstract these sets of features, either to better understand the expert's tacit knowledge or to provide cues and guidance for decision-making that could echo expert performance. In this sense, the ambition would be to design user interfaces that can enable "skill-based" activity. Beyond this, however, is the argument that "skill" is a far broader concept than Rasmussen seems to accept.

Defining Information to Support "Skill-Based Activity"

From the previous discussion, we could say that "skill-based" action involves recognizing and responding to affordances within the environment. In this respect, the word "skill" is intended to capture the accumulated experience that allows activity to be "automatic," or performed with little or no conscious awareness. It is the contention of EID practitioners that traditional user interface designs are created on the premise that all information is of equal value and that the person is able to extract the appropriate information for whatever decision they are making. Given the complexity of process control rooms, this can be challenging. EID[16] responds to a task ecology through the question of what information is necessary and sufficient to enable a person to contribute to the system achieving its functional purpose. From the work domain analysis (figure 5.3), one can determine the relationship between functions at different levels, and from the decision-ladder (figure 5.4), one can determine the strategies that people are liable to follow to perform these tasks, and hence, what information is required to support these tasks. More than this, the abstraction-decomposition space indicates which "values and priorities" constrain system activity and so indicates where information needs to be provided to show the status of the system in terms of these constraints.

EID seeks to reveal the constraints and relations in complex systems. In this way, the task ecology in which a decision maker operates is made clear, so that effort can be directed to making decisions rather than collating information in order to determine system state. EID is intended to support the user in diagnosis of problems through means-end problem-solving strategies. Additionally, the visualization of the system's activity in terms of these constraints and relations means that the operator will be able to respond to a set of features presented as patterns (i.e., the creation of a Gestalt of related information and also in terms of the response that users are expected to perform), relying on a perception-action coupling between what is seen and what action can be performed. In a sense, EID has the aim of "keeping related things together."[17] For the experienced user, the patterns presented by the user interface can be responded to in an automatic, skill-based manner.

The primary goal of the designer, in this instance, is to create a user interface that supports "skill-based" action because that aligns with the perception-action coupling that underpins embodied cognition. In other words, a well-designed display ought to support “at-a-glance” reading and “intuitive” response (both of which are recognized as attributes of well-designed displays in the HCI literature). While this is a logical conclusion to be drawn from this distinction, I wonder whether it is entirely consistent with the broad aim of EID. That is, the knowledge- and rule-based approaches seem to me to rely on information-processing notions of cognition, while the skill-based approach does not. So, either we need to produce separate designs, drawing on distinct theoretical traditions, or we seek to reconcile these into a single approach. My preference is to see whether we can continue with the embodied cognition approach and ground EID in this.

Ecological Interface Design

The foundations of EID involve a mix of concepts from cognitive psychology (specifically from information processing and the study of fault-finding and problem-solving), Gibson’s notion of ecology, systems engineering, and cybernetics. From the latter, Ashby’s “Law of Requisite Variety” would mean that a complex system needs a complex controller, meaning that the combination of human operator, sensors, and automated analysis of data need to be sufficient to mirror the process. The purpose of the user interface would be to expose the human operator to the complexity of the process in such a way as to enable effective control actions to be performed. Further, because a controller requires a model of the system that it is controlling, the operator must be able to access a “model” of the process. From the information-processing approach, such a model would involve the human extracting information-as-content and constructing a “mental model” to make sense of the process. From an embodied perspective, the human operator is considered to be part of the system and, as a consequence, the essential concern is to define the constraints under which the system behaves. We noted, in chapter 4, that Gibson regarded such constraints in terms of the physical laws that governed processes and the conventions (i.e., social, legal, economic, and so on) by which activity is permitted. The abstraction-decomposition space represents these combinations of constraints. Consequently, user interface design should reflect the constraints in the work domain in which the system operates. Fundamentally, this criterion raises the questions of how these constraints can be presented in a relevant manner for the user and what information is required for such presentation.[18]

Thus, the primary difference between EID and "conventional" user interface design is that the latter presents information-as-content, allowing the user to apply different rules of interpretation to extract information into salient combinations, while the former presents information-as-context pertaining to specific purposes.

The essence of EID is the visual representation of control-relevant relations. This reflects the genesis of the approach in process-control domains. Broadly, the overarching aims of the approach are to (a) reflect the task ecology as experienced by the people who work in that domain and (b) present sufficient information to support activity at the different levels of the abstract-decomposition space. For Rasmussen,

> Any control action activated through a work station, serves to change the internal, causal or intentional constraints to let them bring system state to the intended target. The interface should then represent the actual state of affairs in the work space in a way comparable to a representation of the intended or the useful state defined by the current goal, together with the situation-dependent "affordances" i.e., the options available for action on the constraints of the internal processes as defined by the physical design or on intentionality as defined by policies, practices, or regulations.[19]

In order to achieve these aims, Vicente and Rasmussen proposed two design principles:

1. For "skill-based behaviour": To support interaction via time-space signals, the operator should be able to act directly on the display, and the structure of the displayed information should be isomorphic to the part-whole structure of movements;
2. For "rule-based behaviour": Provide a consistent one-to-one mapping between the work domain constraints and the cues or signs provided by the interface.[20]

These principles derive from the aims of EID and can be elaborated further into concepts that inform design decisions.[21] I will illustrate these concepts with examples from the DUal REservoir System Simulation (DURESS)[22] simulator which is one of the best-researched examples of EID. The "mimic concept" draws on the idea (common in process-control user interfaces) of presenting the layout of components in a plant in terms of their structural relations, for instance, in terms of diagrams that show how components connect to each other.

Where EID differs from "conventional" user interfaces is that, in addition to the physical relations between components, there is an emphasis on functional relations. In figure 5.5, a series of valves (indicated by the black

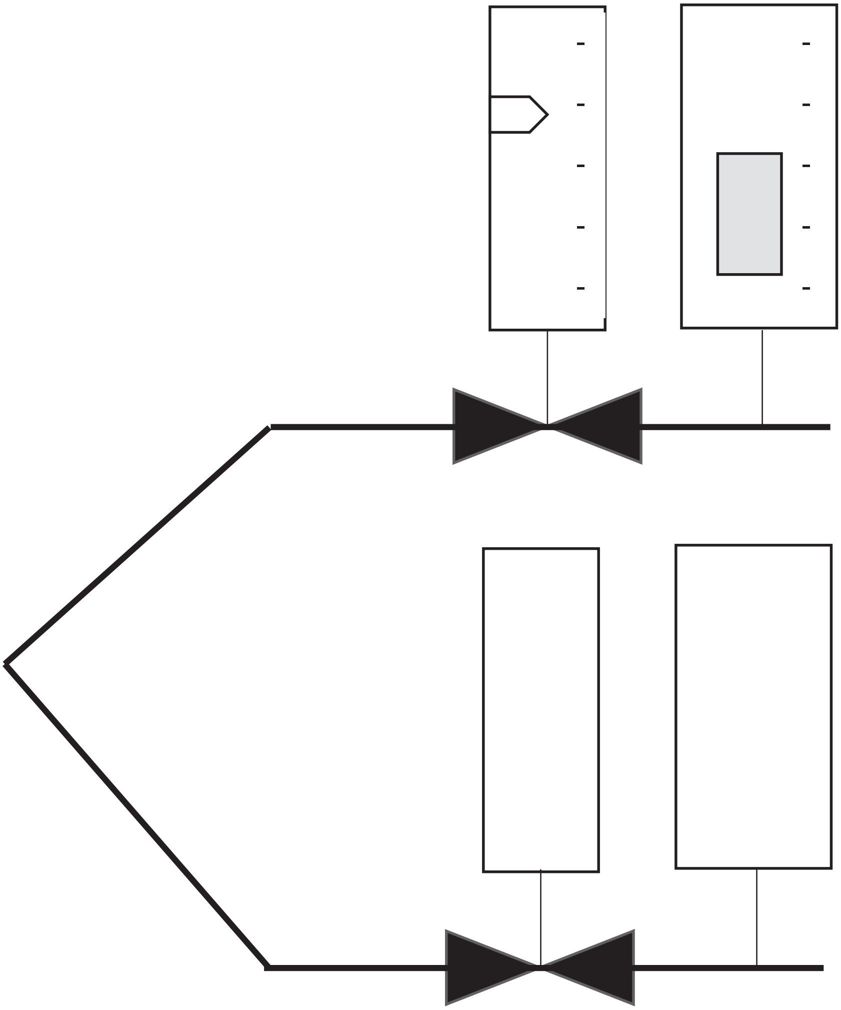

Figure 5.5
Indicating valve state and flow rates in an ecological interface.

“bow-ties”) are connected to control the flow of liquid in a plant. That is, the user interface would show not just the “static” view of the system being controlled but how it changes state. In this way, the user interface will visualize the “flow geometries” that are relevant to the process (and are much like Gibson’s “safe field of travel” for driving or walking, discussed in chapter 3). The “equality principle” states that some parameters describing the system ought to be equal under normal operations (and if these parameters deviate from these levels, then this indicates that the system is out of normal). In figure 5.5,[23] the status of each valve in indicated by a graph showing the flow rate (using a pointer on a scale), together with indicators of flow.

The "conservation principle" states that some parameters will be balanced under normal operations—that is, that energy lost in one part of the system ought to balance the energy gained in another part. One of the more commonly used examples in EID is the mass-energy balance display. In figure 5.6,[24] the quantity of material flowing into the process ("mass in") relates to the quantity of material flowing out. The values of these parameters are shown by the vertical bars, with a connecting line to draw attention to the

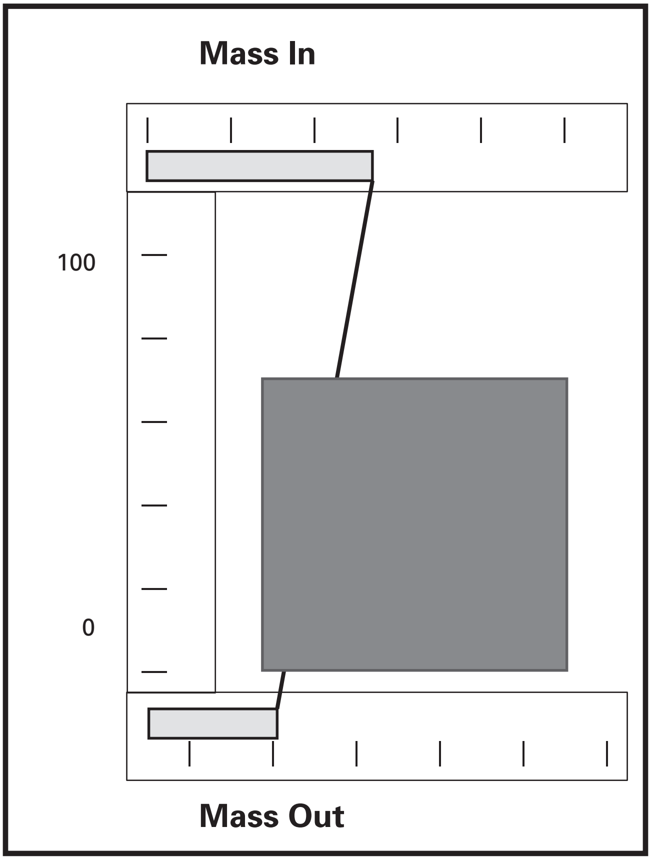

Figure 5.6
Mass-balance display in an ecological interface.

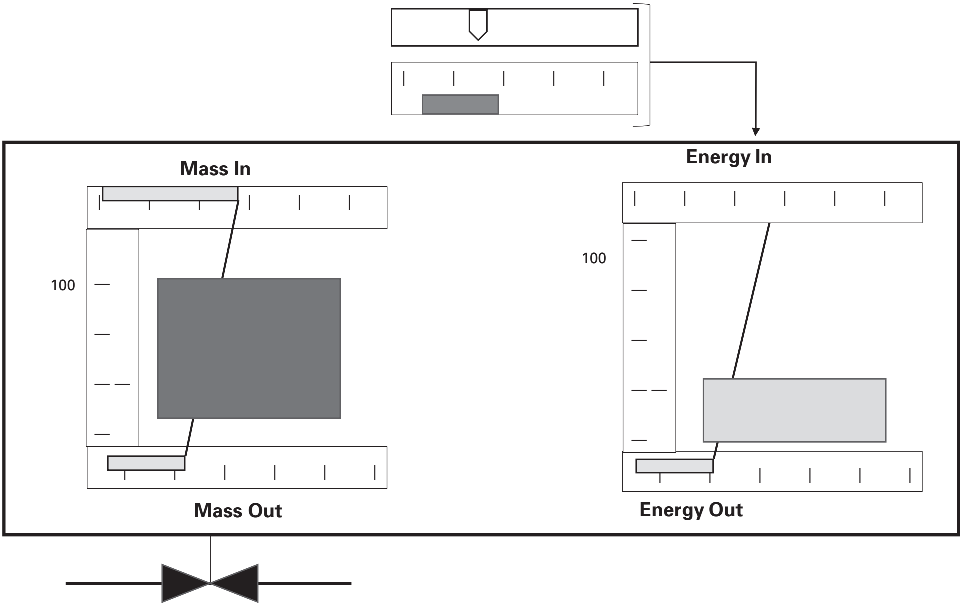

Figure 5.7
Graphically defining relations in an ecological interface.

relative difference within safe limits, defined by the square. In this way, the visual appearance of information creates a Gestalt, or pattern, that can be apprehended as indicating a particular system state (or deviation from that state).

Finally, the "pairing concept" states that certain parameters can be paired to produce a two-dimensional graph, as for temperature flow rate, for example (figure 5.7).[25]

Combining these different principles and concepts results in a composite display in which the elements provide different views of the "system" in terms of the abstraction-decomposition space. For example, figure 5.1 shows a concept user interface for military course of action planning.[26]

Are Ecological Interface Designs Better than Traditional Designs?

Gibson saw computer displays as providing "mediated indirect knowledge at second hand."[27] This would imply that their purpose could only be to display information-as-content, which the user would have to read and understand. The EID approach seeks to present information-as-context. In complex or nonroutine tasks, user performance (measured in diagnostic accuracy) tends to be superior for EID compared with "conventional" user interfaces[28] (although there is evidence that, even for normal operations, EID can be advantageous[29]). One explanation for this advantage is that EID presents information that is salient to understanding the function of the system in that situation while the "conventional" displays might require users to extract and combine information. The part-whole and means-ends approach to design means that the "window" of the process can be adapted to specific functions, which can enhance performance.[30] Indeed, adding information by moving up the abstraction-decomposition space, as the task becomes more complex, can enhance performance. In terms of decision time, some studies suggest that EID can be faster than "conventional" user interfaces, while other studies suggest that it can be slower. One suggestion is that this difference in time can be explained by differences in visual search patterns. However, some "experienced" operators are skeptical about the value of EID (even when using these designs produces measurable improvements in their performance) because they have a repertoire of interpretative strategies through which they define salience. When the "skill" of the user has developed to a level where they can "automatically" see patterns in their environment (or, at least, in the computer displays in their environment), they have developed perception-action coupling that allows them to determine the state of the process. The patterns presented by EID, in such cases, compete with these learned patterns. The question is whether the experienced users can, with practice, replace their old patterns with these new ones. One would expect the answer to be yes, but also that the "old" patterns, as manifestations of tacit knowledge, might be less easy to define that the EID patterns. That is, EID should allow users to "to better judge the validity and quality of specific computer advice,"[31]

What Does Ecological Interface Design Tell Us about Radical Embodied Cognitive Science (and Vice Versa)?

While the EID approach, with its acknowledgement of Gibson's contributions to "ecological psychology" and its version of "affordance," could share some common ancestry with RECS, I would like to argue for a stronger relationship. In part doing so calls for a separation of EID from the information-processing approach that CWA inherits from Rasmussen's early work on fault-finding in electronic circuits. In this early work (and the way in which the sequential decision-making seems to be mapped onto management of complex processes), there are some assumptions that seem to contradict, or reduce the importance of, the "ecological" stance that the approach offers. This is, I feel, compounded by the arbitrary division of "cognition" into "skill-," "rule-" and "knowledge-based behaviors" Vicente and Rasmussen noted that "because operators may engage in higher levels of cognitive control (e.g., knowledge-based action) even if the interface is designed to encourage lower levels, merely supporting the lower levels is not sufficient."[32] How are we to read this statement, if we want to argue against an "information-processing" approach (or, indeed, the notion of levels of cognitive control?). In the first place, this is advising against an overly simplistic assumption that there could be different forms of user interface that directly and uniquely correspond to the levels of control. That is, it does not make sense to design an ecological interface that is specifically for skill-based behavior and to have this as different in kind to one designed for rule-based or knowledge-based behaviors for the simple reason that people will approach the user interfaces with different task constraints and for different decision-making requirements. Consequently, defining the display specifically on the assumption of skill-based behavior could, paradoxically violate the broad aim of EID by presenting information that is either not salient to the task at hand or that requires interpretation in order to make sense of it.

In addition to the suggestion that it is a mistake to design EID explicitly for each "level," I am also proposing that the very idea of "levels" (skill-based, ruled-based, knowledge-based) is misleading and that it makes more sense to approach the design from the single perspective of RECS. First, both RECS and EID are concerned with understanding the relationships between constraints (actor, task, or environment) within an ecological niche. This means that the salience of information is defined by its utility in defining patterns in the environment that accord with actions that a given actor can perform. Different actors (defined by capability or by experience of that environment) could have different definitions of salience. Thus, it makes more sense to define information-as-context in terms of these definitions of salience than in terms of general notions of levels of cognitive control. That is, the EID for someone learning to control a process would differ from that required by an experienced operator—*not*, I think, because the "learner" operates on a knowledge- or rule-based level and the "experienced operator" on skill-based one, but because there are essential differences in kind between the ecological niches of these people. This relates to Brunswik's notion of ecological validity in terms of salient information.

Both EID and RECS are concerned with the notion of "affordance." As I argued in chapter 4, affordance is more than perception-action coupling (and definitely not some property of objects). As such, the notion is antithetical to an information-processing approach; you can have one or the other but not both because affordance emphasizes information-as-context and information-processing emphasizes information-as-content. However, the contrast between information-as-content and information-as-context raises a larger quandary in terms of what it is that the designer (of a user interface in this instance) is designing. If user interface design is a matter of providing information-as-content (on the assumption that users will apply their knowledge to interpreting, and discovering relations within, the information), then the key activities revolve around the definition of content. Indeed, in the mimic displays of process control rooms, each sensor (in the plant) was accorded its own "display" (often a dial on the wall-panel) positioned adjacent to the image of the element that it was sensor of. So, as in figure 5.7, each value would have a display to show the flow-rate through them (and the images of the valves themselves would probably be illuminated to indicate whether or not they were open). In this way, looking at the mimic display (or pacing along the length of the display, in a larger control room) would allow the operator to determine the state of the elements in the process. What this does not tell you, of course, is the relation between these elements (without significant effort on the part of the experienced operator).