Searching for Meaning in Complex Databases: An Ecological Perspective

J.M. Flach and K.B. Bennett

Psychology Department

Wright State University

P.J. Stappers and D.P. Saakes

Industrial Design Engineering

Delft University of Technology

> *One fundamental challenge for the design of the interactive systems of the future is to invent and design environments as open systems in which humans can express themselves and engage in personally meaningful activities.* Fischer (2000, p. 283)

As reflected in the opening quote, a general premise of this chapter is that a target in the design of any technology, such as the web (or other complex databases), is to enhance the "meaningfulness" of the human experience. First, we will introduce the concept of "ecology" in order to explore the source of meaning and the relation of meaning to human experience. Second, we will provide an overview of Cognitive Systems Engineering (CSE). CSE is a relatively new approach to work analysis that explicitly tries to describe the ecology of meaning and the implications for strategies that cognitive agents might utilize to deal with meaning. In Part 3 of the chapter we will focus on the Web as a particular work domain, to consider the special dimensions of this domain that might guide generalizations from previous research on human-machine systems. Part 4 will present two examples. The first example is the Bookhouse interface (e.g., Pejtersen, 1992). The second example is the Multidimensional Scaling Interface developed by Stappers. A concluding section will try to summarize an ecological approach to web design and point to potentially interesting directions for research and design.

What is an ecology?

> *The notion of affordance implies a new theory of meaning and a new way of bridging the gap between mind and matter. To say that an affordance is meaningful is not to say that it is "mental." To say that it is "physical" is not to imply that it is meaningless. The dualism of mental vs. physical ceases to be compulsory. One does not have to believe in a separate realm of mind to speak of meaning, and one does not have to embrace materialism to recognize the necessity of physical stimuli for perception. (Gibson, 1972/1982, p. 409).*

It is impossible to consider any cognitive system (whether human information processor, artificial intelligent agent, or a multi-agent human-machine system like the web) without facing the issue of meaning. What is meaning? Where does it come from? Is it a property of the physical world? Or is it a product of cognitive processing? On one hand, it could be argued that the world presents physical consequences that are independent of the judgments of any cognitive agent (e.g., the consequences of a high-speed collision). Thus, the meaning of an event might be specified independent of any processing or judgment on the part of a cognitive agent. The collision will have the same consequences independent of what an observer might think or believe. On the other hand, two observers might interpret an event differently depending on individual differences, such as differing intentions (e.g., the predator may be trying to create a collision, whereas the prey may be trying to avoid collisions). In this case, the meaning of the collision might be very different, depending on the point of view (prey or predator).

Putting this in the context of information system design, designers who associate meaning with the physical world are likely to seek “objective” standards of meaning. For example, are all the objects in the inventory correctly classified in the web database? Where “correctly” is measured relative to some objective normative criterion (e.g., Dewey decimal system or other convention for cataloging inventory). The focus of design efforts would be on the correspondence between the representations in the Web and the physical objects (e.g., books or roller skates) that are being represented. Typically the degree of correspondence would be measured against some model of the domain that specifies both what belongs in the database and what differences make a difference (i.e., the dimensions of the space). This model will often reflect some degree of intentionality, but this will normally reflect the intentions of professionals within a domain (e.g., librarians), rather than general users (e.g., library patrons looking for an interesting book).

Designers who consider meaning to be the product of cognitive processing are likely to look to the human users for “subjective” aspects of meaning. What is important (meaningful) to the human users? Designers who take this perspective are likely to be interested in the “mental models” of their users. They will typically measure their designs in terms of their correspondence or match to the users’ mental models. How do the users think about the objects in the database? Naturally, this does not require that attributes of the physical world be ignored, but they become secondary considerations. Designers who focus exclusively on what the users’ think have no basis for differentiating between perspicacity and superstition. Does the mental model reflect a well tuned or a naïve map to the actual work domain?

An ecological approach offers a middle path as illustrated in Figure 1. The term “ecology” is used to reflect a relation between animal and environment, similar to the construct of “Umwelt” used by von Uexküll (1957). It is important to note that this relation is not simply the sum of the two components (physical world + animal), but the “fit” or “relation” between the two. The ecology reflects emergent properties of the interaction of an animal with an environment. As such, it has a unique dimensionality that is different from those of either the animal or the environment. Figure 1 illustrates critical dimensions that emerge from this relation between animal and environment.

The three dimensions of the ecology that are important for assessing the “fit” between animal and environment are represented by the verbs “satisfying,” “affording,” and “specifying.” “Satisfying” represents the value constraints on the interaction. This dimension reflects the intentions or goals of the animal in relation to the consequences of the environment. In the context of web design, satisfying reflects both the goals of the user searching the web and the objective properties of the websites or the products represented by these sites. Do the “hits” from a search match the information needs of the person doing the search? Does the product selected through a website satisfy the needs of the user? From the perspective of design, this reflects a need for the designer to have insights both about what people want and about what can be offered?

“Affording” represents the action constraints on the interaction. It reflects the effectivities of the animal in relation to the possibilities offered by the environment. This represents what the animal can do. “Specifying” represents the information constraints on the interaction. It reflects the perspicacity (perceptual skill) of the animal in relation to the structure (e.g., optical flow) in the environment. The confluence of “satisfying,” “affording,” and “specifying” create the possibility for control. That is, “satisfying” reflects the reference signal, “specifying” reflects feedback, and “affording” reflects the action to reduce discrepancies between the reference and feedback information.

In the context of database design, the “possibilities” represent the information and products available through the database and the “structure” would represent the measurable distinctions that might be used to organize or classify these possibilities. The “effectivities” and “perspicacity” of the user will largely be determined by the user interface. That is, the interface provides both the means for action and the feedback about the opportunities for and consequences of actions. The kinds of effectivities typically provided by the web include pointing and clicking on graphical objects, selecting items from menus, and the Boolean logic used to initiate a search. The feedback includes graphical and text descriptions of data, products, and options. For the design to function as a control system the feedback must be comparable to a reference (reflecting the intentions of the users and consequences of action) and the comparison must result in a clear specification of the actions necessary to reduce any discrepancy (i.e., error) from the reference.

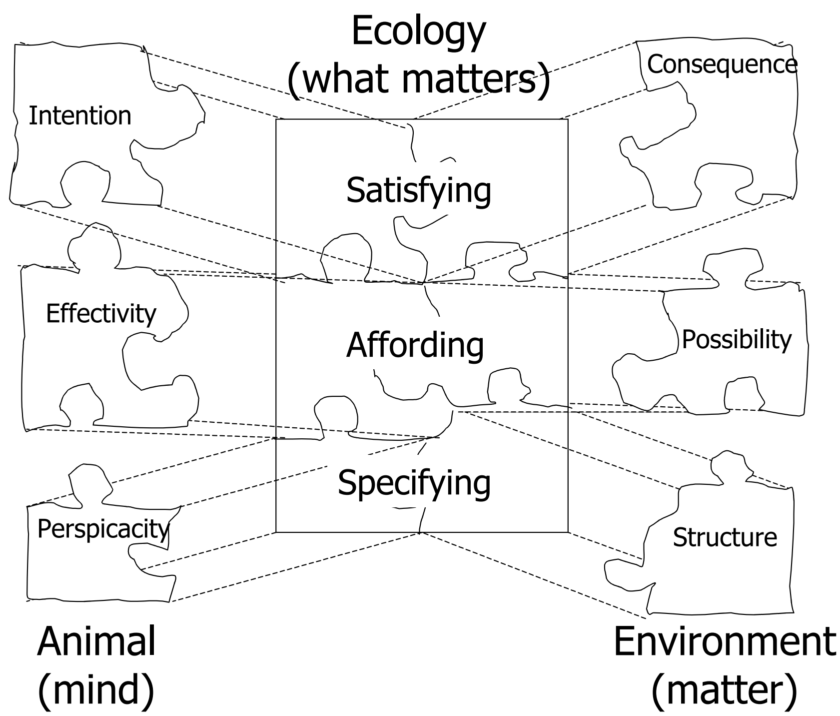

Figure 1. The ecology is a set of emergent properties resulting from the fit between an animal and an environment. Some of these properties are represented by the verbs "satisfying," "affording," and "specifying."

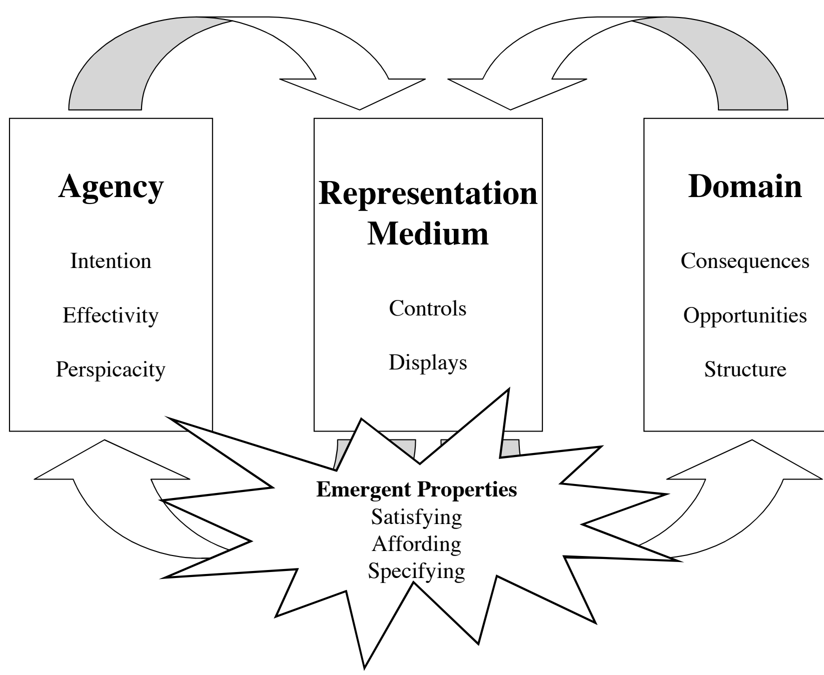

Figure 2. The dynamics of perception action coupling. The coupling of performance (control) with exploration (hypothesis testing) to form an adaptive control system.

The choice of verbs to specify key dimensions of the ecology was intended to emphasize the fact that the ecology is not static. The ecology reflects a dynamic coupling of perception and action. Figure 2 was designed to emphasize this dynamic. The dynamic is depicted as two coupled loops (a "Figure 8"). Movement from the top left to the lower right then looping back to the center represents the control problem where the center box (representation medium) functions as a comparator in which intentions and consequences are compared to produce an "error" signal which drives actions to reduce the error (satisfy the intention). This is essentially the "cybernetic" image of a cognitive system.

The backward loop from the domain (e.g., consequences, opportunities) through the representation to an agent represents an abductive approach to the semiotic problem. That is, how the meaning of objects and events in the domain are "mapped" (i.e., communicated) through a representation (e.g., symbol, sign, or signal) to a belief system. This loop reflects the observation process that allows the construction of expectations and hypotheses about the world. These beliefs can be tested, tuned, or otherwise improved by exploratory actions on the domain.

One of the key insights of an ecological approach is that the control and abductive processes are dynamically linked in such a way that neither is logically prior to another. This is in contrast to classical information processing models where the semiotic problem (perception) is treated as logically prior to control (decisions and action). It could be argued that most approaches to cognition either ignore one aspect of this coupling (e.g., focus on behavior/activity without any theory of meaning; or focus on symbol processing (semiotics) without any theory of action) or they parse the cognitive process in a way that the control and semiotic problems are independent. On the other hand, an ecological approach is predicated on the assumption that the dynamic coupling is the heart of the problem and that understanding the emergent properties of this coupling should be the goal of cognitive science. This also reflects a vision of "design" where problem and solution, question and answer, evolve in mutual enhancement, not in succession (e.g., Schön, 1983).

In fact, the case could be made that all of the dimensions specified in Figures 1 and 2 (those attributed to the animal/agent, to the environment/domain, to the representation, and to the ecology) are emergent properties. This is consistent with William James' (1909) philosophy of radical empiricism, that "meaning" can only be measured relative to its impact on experience. This is in contrast to "realist" or "idealist" positions that posit extrinsic standards for meaning in either an objectively "real" world or some abstract, logical "ideal." This radical empiricist view seems to be most consistent with our observations of web use. For example, often a user will begin a web search with only an ambiguously specified intention. However, this intention often is refined in the process of interacting with the web – to the point where the satisfactory end of the search has little resemblance to the intention that initiated the interaction. Thus, the intention itself is an emergent property of the dynamic, not an extrinsic fact imposed on the dynamic. This is one of the key challenges of web design – to help people to satisfy their intentions – even when these intentions are ill-formed or impossible to articulate in advance. The users often won't be able to articulate a clear goal – even when they can reliably recognize a solution (when and if they get there). So, the key point here is that cognitive processes can not be reduced to ordered sequences where intentions are refined into programs or plans that in turn direct actions. Cognitive systems are dynamic! They learn by doing – by trial and error (in Piagetian terms, by assimilation and accommodation). They must project their experiences into an uncertain future (assimilation) and they must revise their beliefs and expectations based on the experiences that result (accommodation).

The dynamic coupling of perception and action illustrated in Figure 2 is our conceptual image of "experience." The emergent properties of satisfying, affording, and specifying are dimensions of meaning. Thus, meaning is neither an "output" (or product) of processing, nor is it an "input" to processing. But rather, it is a property of the dynamic, similar to "stability." When the representation provides rich structure that matches well with the consequences and opportunities in the domain and the effectivities and intentions of an agent, then the dynamic can be expected to stabilize in a fashion that is "satisfying." This is what we mean by "skilled" interaction, and the sense is one of direct access to meaning – the experience is "meaningful." When the representation does not match with these dimensions, then instability will be the likely result. Or perhaps, stability in a fashion that is "unsatisfying." The sense is one of confusion or "meaninglessness."

Cognitive Systems Engineering

The problem we face in modeling systems incorporating human actors is, however, that humans do not have stable input-output characteristics that can be studied in isolation. When the system is put to work, the human elements change their characteristics; they adapt to the functional characteristics of the working system, and they modify system characteristics to serve their particular needs and preferences. Rasmussen, Pejtersen, & Goodstein, 1994, p. 6)

Cognitive Systems Engineering (CSE) is motivated by the simple question: What can designers do to ensure or at least increase the probability that the cognitive system will stabilize in a way that satisfies the functional goals motivating the design. This question is particularly difficult in today's environment where the work domain is changing at a rapid pace. These changes reflect the opportunities of evolving technologies such as the Internet and the competitive demands for innovation that result. When the pace of change is slow, stability can often be achieved by standardization around well-established procedures. But when the pace of change is fast, operators are often called on to "complete" the design by adapting to contingencies for which no well-established procedures exist.

In a sense, the goal of CSE is to design "adaptive" systems. In this context, the "human factor" takes on new and interesting dimensions. Classically, the human was viewed as a collection of cognitive limitations – a source of error. Designers tended to focus on ways to minimize these errors – to ensure that the humans did not deviate from the well-established procedures. Today, however, focus is shifting to the "adaptive" and "creative" abilities that humans bring to the work place. The focus of CSE tends to be on ways to leverage these abilities against the uncertainties of a complex and rapidly changing work environment. The focus is on providing the humans with the tools they need to invent new procedures as demanded by the changing work context. Figure 3 shows three fulcrum points that designers might use in leveraging the adaptive abilities of humans against the demands of complex work environments: work analysis, ecological interface design, and instruction design.

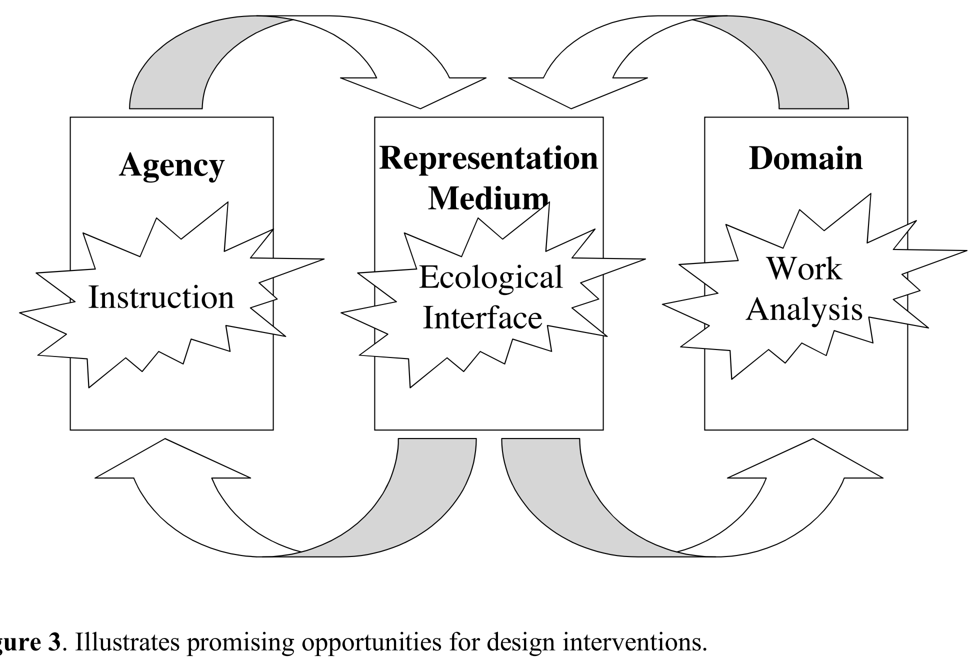

Figure 3. Illustrates promising opportunities for design interventions.

Work Analysis

Ironically, the first place to start, in preparing to deal with change, is to search for regularities (lawfulness, invariance, or constraints). The regularities become the "landmarks" that can guide exploration. Classically, primary sources of regularity were the "standard procedures." Thus, classical work analysis focused on the "activity" of work, often in relation to "one best way." However, in an environment where procedures must be adapted to changing work demands – the focus must shift from the procedures to the work demands themselves. In the terms of Simon's (1981) famous ant analogy, the focus shifts from the path of the ant to the landscape of the beach (see Vicente, 1999).

Flach, Jacques, Patrick, Amelink, Van Passen & Mulder (In press) have made an analogy to the distinction between "survey" and "route" knowledge in the navigation literature. Task analysis describes the work domain in terms of routes. That is, sequences of actions or procedures. For example, "go to the menu and select the third option, then a dialogue box will appear, fill in the information in the dialogue box and then click on the continue button, etc." Work analysis attempts to provide a more complete map or "survey" of the workspace. For example, a tree diagram showing the complete menu structure making all options explicit.

Route knowledge has several attractive features relative to computers. First, it is easily encoded in a rule-based description. Thus, it fits with the GOFAI (good old fashion artificial intelligence) image of information processing and it is easily conveyed via text-based or verbal instructions. Second, it tends to reduce memory load (the number of rules) by focusing only on the optimal paths (e.g., minimum number of steps) to accomplish the intended functions; and at the same time it naturally inhibits deviations from the optimal paths (which may be especially important for safety critical systems). The major drawback of route knowledge is that it tends to be brittle. That is, it works only as long as the environment is static or stable. However, if the environment changes so that a prescribed route is blocked, if destinations/goals other than those specified in the formulation of the routes become desirable, or if unexpected events force an agent from the prescribed route, then route knowledge provides little support for adaptation.

Survey knowledge tends to require a higher degree of integration. That is, unlike route knowledge, which is nearly a direct analog to activity, survey knowledge requires the distillation of abstract relations (e.g., cardinal directions) from the activities. Survey knowledge is difficult to capture with rule-based descriptions or to communicate verbally. It typically requires a spatial form of representation (e.g., map, tree diagram, flow chart). The advantage of survey knowledge, however, is that it tends to be robust. That is, it provides a foundation for generating new solutions (e.g., a new path), when the environment changes (e.g., the normative route is blocked or a new goal emerges).

CSE was developed to address the demands for adaptation in complex work environments like the web. Change is a universal in these environments. In such environments rule-based or route descriptions will have limited utility. These environments demand the extra effort required by survey representations. Thus, work analysis represents an attempt to survey the work landscape. And CSE is a belief that design decisions should reflect "survey" knowledge of the workspace.

There are two useful conceptual tools for surveying complex problem spaces: decomposition and abstraction. Decomposition attacks complexity by partitioning it into smaller chunks. The hope is that by addressing each chunk in turn we take small steps toward understanding the larger problem. The trick is to find a partitioning that preserves the meaning of the whole. That is, you hope that the steps will "add up" to a deeper understanding of the domain. There are many ways to "break up" a complex problem, but few ways that will preserve the integrity of the whole.

Abstraction attacks complexity by introducing conceptual structures. For example, the "gist" of a story could be preserved even though all the parts (words) were changed. In this sense, the "gist" is an abstraction. Another example is categorization. A category such as "furniture" reflects a relation among objects that cannot be reduced to a list of common parts. Abstractions generally function to highlight dimensions that are thought to be "essential," while blurring other dimensions that are considered less important. Abstraction is illustrated very clearly in Hutchins (1995) discussion of different forms of maps. Maps are not simple iconic images of the layout of the earth. Rather they are carefully crafted abstractions designed to highlight some dimensions (e.g., line of sight directions) at the cost of blurring others (e.g., area).

Rasmussen (1986) reminds us that we need both strategies and that the strategies are not independent. The dimensions highlighted and blurred by the abstraction process have implications for what decompositions will be appropriate. Figure 4 attempts to illustrate this interaction between abstraction and decomposition. Abstraction is illustrated as the vertical dimension of this space and level of decomposition is illustrated as the horizontal dimension.

Although five levels of abstraction are represented in Figure 4, consistent with Rasmussen's (e.g., 1986) Abstraction Hierarchy, we will not attempt to define these levels or to make any claim about the absolute value of these particular levels. The claim that we do want to make here is that no single level of abstraction provides a comprehensive survey of the meaning landscape. Abstractions that highlight the dimensions associated with satisfying (e.g., intentions and consequences) often will blur the dimensions associated with affording or specifying and vice versa. The diagonal regions of the matrix highlighted in Figure 4 reflect experiences from previous work analyses and from protocols of experts during problem solving in their domain (Rasmussen, 1986; Rasmussen et al. 1994; Vicente, 1999). These experiences suggest that questions associated with "why" (satisfying) are typically best seen with high degrees of abstraction and relatively gross levels of decomposition. This reflects a "top-down" view of work. Questions associated with "how" (affording and specifying) typically require more concrete descriptions and finer levels of decomposition. This reflects a "bottom-up" view of work. The "bottom-up" view helps to answer questions about how to carry out a particular procedure. The "top-down" view, however, becomes critical for deciding which of two procedures might be best in a given context.

Solving the adaptive control problem typically requires iterating between the "bottom-up" view and the "top-down" view. The bottom-up view helps to clarify the best way to carry out a particular strategy. The top-down view is useful for evaluating different strategies to assess which of several strategies might be most appropriate for a given context. Historically, the entry of CSE into many work domains has been a controversy over different possible strategies or over the appropriate criterion for switching from one strategy to another. This reflects the realization that there may be more than one "best way" and that the best way may vary with the changing context.

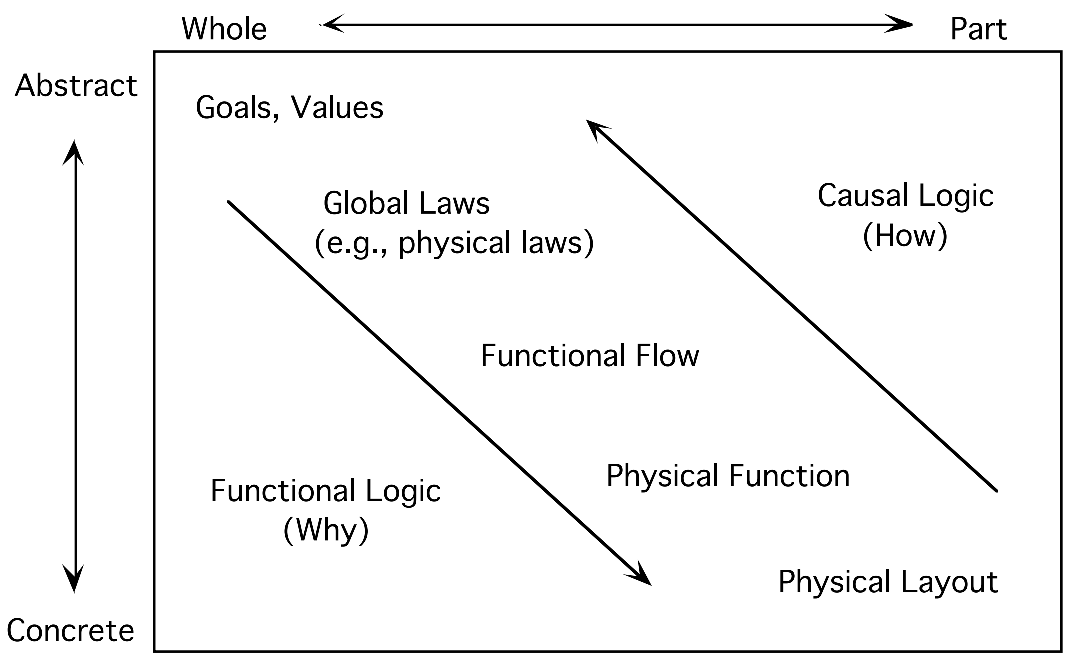

Figure 4. Interactions between abstraction and decomposition for attempts to understand means (how) and ends (why) within a complex work domain.

In sum, an important step in helping humans to respond in an adaptive way to meaningful properties of a work domain is to identify those properties. Fischer (2000) refers to this as "domain-oriented design." The goal is to "support not only human-computer interaction, but human problem-domain communication, making the computer invisible and bringing tasks to the forefront" (p. 284). Work domain analysis is a process for surveying the work landscape to discover meaningful constraints. A comprehensive survey requires multiple perspectives in terms of the levels of abstraction and the degree of decomposition.

A final word of caution: understanding of a work domain is an asymptotic process. That is, it is never complete. Even for a simple system, such as a telephone, design does not end with the specification of the product features. This is illustrated by a public phone in the Netherlands that featured a coin slot positioned at the top of the phone, just above the hook holding the horn of the phone. Unfortunately, in many places the phones were hung at eye-height (perhaps to make reading the numbers easy). However, the coin slot was now occluded from sight so that many people could not "see" where to put the money – and thus could not operate the phone (as reflected on handwritten signs glued to the phones or mounted on the walls). The point is that the people who hung the phones and the people who wrote the notes were participants in an evolving design. Thus, for more complex work domains there will always be new facets to discover. Work domain analysis is not something that is completed prior to design. It is an iterative component of an ongoing design process. Each new design is a step in an evolutionary process. It is a "test" of the current survey knowledge. It is important that there are mechanisms so that the survey knowledge can be refined based on the feedback from each test and to allow users to participate in the design process.

Ecological Interface Design

The goal of Ecological Interface Design (EID) is to explicitly represent the meaningful constraints that have been discovered through work domain analysis in the interface. An assumption of this approach is that agents who have an explicit survey representation (e.g., map) of the work landscape will be able to adapt their choice of routes to consistently satisfy the functional goals in the face of a changing world. Thus, the "push" for ecological interfaces comes from the increasingly fast pace of change in the work place and the demand for creative adaptations that results. The "pull" for ecological interfaces comes from the flexibility of graphical interfaces. As we noted in the discussion of survey knowledge above – it is difficult to communicate survey knowledge in the form of rules or verbal instructions. Typically, some form of spatial representation is required. Graphical interfaces provide the opportunity to build these representations.

The idea of EID is tightly coupled to the CSE approach to work analysis. The general notion is to make the nested hierarchy of constraints that has been identified through work analysis explicit to the operator. This has typically involved graphical representations using configural representations as illustrated in Figure 5. In some sense, the structure in the interface should mirror the structure of the work domain. Thus, the interface might have a geometry in which global properties reflect higher levels within the abstraction hierarchy (e.g. functional goal and physical laws). Nested within these global properties might be regional structures that reflect the organizational and physical relations within the work domain. Finally, the geometric details might map to specific states within the work domain. The purpose of this structure is to help the operator to "see" whether they are satisfying the functional goals; to see the logic of the process dynamic; and to see the individual state variables in the context of the functional and process constraints.

The prototype for an EID is Vicente's (1999) DURESS interface. Vicente used the abstraction hierarchy framework to develop a detailed map of the means-ends relations associated with the thermodynamic DURESS process. He then developed a geometric structure to make these relations explicit in the interface representation. For example, he used a funnel metaphor to depict the flow of mass and energy through the system reservoir.

Vicente's interface represents an ideal that provides a valuable example of the ecological interface approach. However, it is an ideal that may be difficult to realize in work domains such as the web – where there are many more degrees of freedom than in the DURESS example. The economic, political, and social constraints that play such an important role in the Internet are far less stable and less understood than the laws of mass and energy balance that constrain a thermodynamic process. While computer graphics and virtual reality systems provide a rich palate for building representations – our understanding of many work domains remains greatly impoverished.

The attribute that distinguishes EID from other approaches to interface design is the role of the abstraction hierarchy for identifying the meaningful dimensions of the experience. However, the approach shares the goals of many other approaches (e.g., Fischer, 2000; Hutchins, 1995; Shneiderman, 1983; Woods, 1984) – that is, to make the interactions with meaningful properties of the work domain as "direct" as possible. Where possible the interface should help the operator to "directly" perceive the work domain – this involves not only seeing what is, but seeing what is coming. Direct perception involves seeing the patterns of unfolding events that specify future outcomes (e.g., whether goals will be satisfied). This ability to anticipate the future is critical to stability in any control system where there are long time-constants between actions and consequences. One way to think about the higher-levels within the abstraction hierarchy is that they represent internal or hidden states of the process. Making these states explicit can be thought of as a form of prediction. For example, the shape of the funnels in the DURESS interface helps to specify the directions and rates of change of the mass and energy inventories. The interface should also help the operator to "directly" manipulate the work domain. This direct manipulation is important both for control and hypothesis testing.

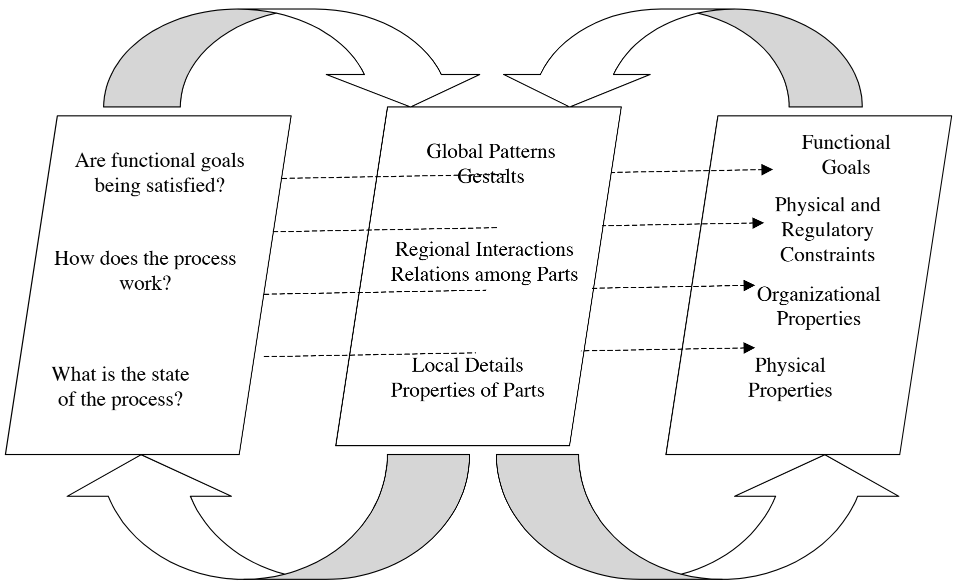

Figure 5. The goal of an ecological interface design is to make the hierarchical means-ends relations in the work domain explicit to the human operators who must navigate among these relations. Note that the items within boxes are assumed to operate in parallel, not in sequence as might be inferred from the arrows entering and leaving the boxes.

Designers who are looking for an easy answer to the interface design problem will be sorely disappointed by the EID approach. Mapping out the abstraction hierarchy is a daunting task – and for domains that are evolving at the pace of the Internet – the meaning landscape is probably changing at a pace that is greater than our ability to keep up. Thus, the mapping problem will be a continuous battle. To the degree that the domain can be mapped – there is still the question of building a representation that captures meaningful aspects of the domain in a way that allows direct perception and direct manipulation. As usual, there are many more ways to do this wrong than right. There is no simple formula that guarantees success. The best that we can do is to provide some examples of promising interfaces in later sections of this chapter. Hopefully, these interfaces will be a source of inspiration. EID is not an answer, but an invitation to engage in the difficult search for meaning.

Instructional Design

A misconception about Ecological Interface Design is that the resulting representation will be "natural." Where "natural" means that people will be able to immediately "see" the deep structure in the representation without instruction or practice, because it somehow taps into native perceptual abilities. Unfortunately, this will rarely be possible. The ecological interface design does not eliminate the complexity or requisite variety that is intrinsic to a work domain. Thus, it does not necessarily eliminate the need for extensive learning. However, it may have important implications for the design of training and instructional systems. By making the constraints explicit in the interface, ecological interfaces support users to work through perception/action, rather than through inference in exploring situations outside the standard operating conditions. This creates the opportunity for "learning by doing"

To the extent that the interface allows direct manipulation and direct perception, it provides explicit feedback that should help agents to discover meaningful aspects of the domain through interaction. When meaning is clear, the interface should allow direct comparison between consequences and intentions in a way that specifies the appropriate action. However, even when meaning is not clear, the interface should allow interrogation through action – so that, the actor can "learn" or "discover" the possibilities offered by the domain. That is, the actor should be able to test hypotheses out at the interface.

Thus, leaning is expected to be an integral part of any adaptive system. Learning becomes part of the job – practitioners are expected to be "reflective" (Schön, 1983). A good interface should facilitate reflection. With an ecological interface the learning process might be more similar to learning to drive a car or to play a sport, rather than learning facts or rules. In this sense, "natural" does not necessarily translate to "easy." Rather, it says that the learning process might be more similar to learning the skill of kicking a soccer ball, rather than memorizing the rules of soccer.

The "Web" as a Domain?

In principle, the web is not a work domain, per se, but rather it is a medium or technology capable of supporting many different types of work. It is what Fischer (2000) has called a High Functionality Application (HFA). He noted that HFAs are complex systems in and of themselves because they serve the needs of large and diverse user populations. The HFA itself has little "meaning." The meaning arises when an HFA is used as a tool within a specific work domain. However, it may be valuable to speculate about the kind of work that might utilize this technology. The Web will most likely be a tool for domains of education, entertainment, and business.

Many of the early concepts of ecological interface, such as Vicente's (1999) DURESS, were designed for the process control industry. In these domains, physical laws (e.g., mass and energy balances) played an important role. However, the web and the information systems upon which it depends are much less constrained by physical laws. Not that the physical laws are suspended – but they become transparent or uninteresting in relation to the functional goals for the work interactions. The structure that will give meaning to work over the Web is more likely to come from social conventions (which we will use broadly to refer to cultural, political, legal, and economic constraints).

In the context of work, the social, political, and economic sources of constraint are typically no less important than the physical laws that govern the flow of mass within a feedwater control system. However, while physical laws tend to be invariant over the history of a process control system, social conventions tend to evolve along with the technologies they constrain (e.g., conventions of point and click operation; etiquette of using cell phones). Furthermore, it is likely that the physical constraints of most process control systems are both simpler (e.g., fewer degrees of freedom) and better understood than the social constraints that govern education, business, and entertainment. This poses a significant challenge for CSE and ecological interface design. The two examples in the following section were chosen specifically to illustrate systems dominated by social conventions – in both cases people must search through complex databases to satisfy personal interests.

Two Case Studies

In the previous sections, we have tried to layout a conceptual framework for an ecological approach to web design. We hope that the conceptual abstractions will provide some insights into the approach. However, we also realize that many details have been glossed over and that there is a wide "gulf of execution" (to borrow from Norman, 1986) between the conceptual abstractions and the details of a concrete design problem. In this section, we will try to bridge this gulf with two case studies that we think have particular relevance for the problem of web design.

Example 1: BookHouse fiction retrieval system

The BookHouse (Pejtersen, 1979; 1984; 1988; 1992) was developed to help general patrons of a library including children find "interesting" books to read. In practice, the graphical interface that allows users to navigate through a virtual library is the most salient aspect of this design. However, from the point of view of CSE the more interesting question is associated with the development of a classification scheme for fiction. Traditional methods for identifying and locating books of fiction in a library are very simple and generally ineffective. Most libraries categorize these books by bibliographic data alone -- author's name, title, ISBN number, etc. Thus, the options for finding a book are rather limited unless the user knows of a specific book and at least some of its bibliographic data. If the user doesn't have the appropriate data or doesn't have a particular book in mind (often the case) the user is left to browse the fiction shelves or to consult a librarian for professional help. The challenge for a CSE approach is to discover meaningful categorical distinctions for classifying fiction. That is, there is a need for a categorical structure that will address the goals of satisfying, affording, and specifying. This challenge is illustrated in Figure 6.

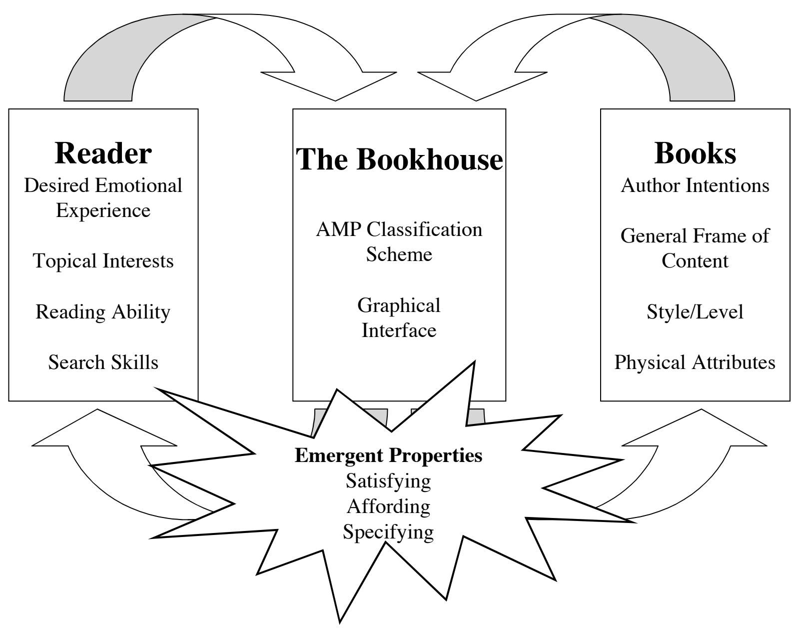

Figure 6. The challenge for the Bookhouse was to develop a classifications system and interface to help a diverse group of users find books that they would enjoy.

<u>Domain</u>. The key domain question was: what attributes of books made a difference with respect to the intentionality of users? To get insights into this problem Pejtersen (1979) examined user queries during actual information retrieval negotiations with librarians. These observations led her to suggest five dimensions of books that reflected users needs: author's intention; literary school, style, and values; frame; subject matter; and accessibility. As a result ot this analysis, five general dimensions have been included in the Bookhouse.

1). The "Bibliographic data" dimension includes the typical information used to classify books, including title, author, date, publisher, etc.

2). The "Subject Matter" dimension includes the "story line" of the book, essentially what the story is about. This includes the topic matter that is being addressed, the elements of the plot, the course of events and the social interactions and relationships that evolve.

3). The "Frame" dimension is essentially a description of the context within which the story occurs. The primary aspects of this dimension are the time (past, present, future), the place (country, geographic location) and the social milieu (mores, issues, reform movements, etc.) that form the setting for the story.

4). The "Author's Intention" dimension captures the reasons why (and to some degree how) the author wrote the book. Authors can write books to provide a variety of emotional experiences: among other reasons authors might strive to amuse, scare, or uplift their audience. Authors might also have ulterior motives: among other reasons authors might strive to inform, sway public opinion, or encite their readers to action. Finally, the manner in which the authors write their books might adhere to particular literary schools, styles, or values.

5). The final dimension is "Accessibility," which refers to a constellation of factors that can influence how well an author's story will be communicated. Some factors are straight-forward factors determining whether or not a reader will be able to read the book, such as the reading level and physical characteristics of the book (e.g., size of the font type). Others will influence the degree to which readers will be intrigued by, or can relate to, the message in the story (e.g., literary form, the main characters and their age).

This type of analysis, aimed at uncovering the sources of regularity in a particular domain, is the foundation of any enterprise that is to be referred to as ecological in nature. Thus, each of the books in the Bookhouse were "indexed" according to how its content fit into the classification scheme; this constitutes the information about that book that will then be available during search. This constitutes the meaning landscape.

In addition to considering the meaning landscape. It is often useful to consider different strategies for navigating that landscape (Pejtersen, 1979). How might people search for interesting books? One strategy is search by analogy. That is, the user specifies a book that was particularly satisfying and asks for a book that is similar. For example, "I want another book like *Gone with the Wind*." Note that "similar" can be problematic, in that everything is similar to everything else in many different ways (e.g. Dreyfus, 1994). Within the Bookhouse, the classification system provides a basis for judging similaritly. Books that have many common dimensions (e.g., similar author intention, similar topic area, similar frame, etc.) will be judge more similar.

Another strategy is an analytic search. With this strategy the patrons specify their desires directly in terms of the classification dimensions – "I want a romance set in the South during the civil war." This strategy depends on the patron knowing the dimensions. An important design issue is who takes the responsibility for the "knowing." Will the patron be required to learn dimensions that are convenient for the database system? Or will the system be designed to match the the natural dimensions used by patrons. For a system like the Bookhouse, where user intentionality is critical to meaning – it makes sense to design the classification system to reflect distinctions that are meaningful to users.

Other strategies for search include empirical search and browsing. Empirical search refers to the situation where a librarian suggests books based on personal attributes of the patron – "Here is the type of book that teenage girls usually enjoy." Browsing could involve picking books out at random or based on cover illustrations (e.g., a woman in period costume in front of a southern mansion) and perusing the contents (e.g., reading the jacket).

Having a description of the meaning landscape and understanding some of the strategies for navigating that landscape provides the foundation for designing interfaces that will help patrons to achieve their functional goals.

<u>Agency</u>. One of the key elements in the design of any successful human-computer system is an explicit consideration of the capabilities, limitations, knoweldge, and other characteristics of those who will use it. Although the admonition to "Know Thy User" has become a platitude for design, it is in fact a very critical

consideration. Field studies and unstructured interviews were conducted. It was found that the targeted user population had a profound lack of homogeneity, capabilities and experiences. The results indicated that the user population ranged widely on a number of dimensions, including familiarity with computer systems, chronological age (and all of the associated differences in capabilities and limitations), familarity with fiction, and reasons for being there in the first place. Two broad subgroups with substantially different sets of capabilities and limitations were identified: children (ages 6 through 16) and adults (ages 16 and up). In particular, the general cognitive development, reading skills, and command of the language were quite different for these two subgroups. This translated into substantially different needs: children and adults tended to formulate their questions about books of fiction using different sets of linguistic terms and different types of semantic content (i.e., at different levels of the AMP classification scheme).

Consideration of these factors resulted in the development of two separate databases of fiction books for use in the BookHouse retrieval system. The children's database contained different types of accessibility information and different descriptions of the books and their content. These changes were aimed at matching the language and the concepts that children used in their natural interactions with librarians. A third database, esentially a combination of the previous two, was also developed to support "inter-database" browsing (e.g., teenagers browsing the adult database or parents browsing the children database).

In summary, the targeted user population of the fiction retrieval system was extremely diverse with substantial differences in a number of dimensions. An additional consideration was that most users were likely to have only casual and infrequent interactions with the system. This stands in sharp contrast to the typical users of complex socio-technical systems (i.e., those systems controlled by the laws of nature -- e.g., process control), who are highly homogoneous in capabalities and have extensive training. It was therefore essential that the system be both intuitive to learn and easy to use. A design strategy was needed that would capitalize upon general conceptual knowledge that was common to all library patrons. More specifically, the system would need to be designed using an interface metaphor that was familiar to this diverse group of patrons.

Figure 7. A "scene" from the Bookhouse. Each person represents a different "mode" for searching the library database.

Ecological Interface. The traditional approach to searching a database is through linguistic commands issued to a computerized intermediary (i.e., keyboard-entered alpha-numeric searches). With the BookHouse system a user completes a search by navigating through a carefully crafted spatial metaphor: a "virtual" library. The user is initially presented with a view from the outside of this virtual library. A fundamental design feature of the system is that visual and verbal feedback is provided to alert the user of the potential for interaction, to specify the action required to initiate that interaction and to outline the goal that the interaction will achieve. In this particular instance the user must enter the BookHouse to initiate a session. The user is informed of this requirement when he/she positions the pointer over the opening of the virtual library (the area framed by the doorway). A white graphical outline of the opening appears (i.e., the icon is highlighted), signifying the potential for action. At the same time a verbal description of the action required (i.e., "Press here") and the end result (i.e., "to get into the Book House") is also provided in a text box located at the bottom of the screen. These conventions (point and click interaction style; visual and verbal feedback on potential interaction, required acts of execution and goals) are used consistently throughout the interface.

Once inside the BookHouse the user is presented with a view of three "virtual rooms" whose opening can be pointed at and clicked on to enter. Each room corresponds to one of the three databases (children, adult, combined) that can be searched. The identity of each database is signified by verbal labels over the rooms, the composition of the figures at the entrance to the room, and the verbal instructions that appear at the bottom of the screen. Entering a room (i.e., pointing and clicking on the opening) is an action that signifies which of the three databases that the user wishes to search.

The user is then presented with a room that portrays various people engaged in four different activities shown in Figure 7. Each of these people-activity icons corresponds to the four different types of strategies that can be executed to search a database; clicking on an icon indicates that the user wishes to initiate that type of search. The four search strategies are referred to as search by analogy, analytical search, browsing through descriptions of books, and browsing through a "picture thesaurus" that graphically represents constellations of books. The analytical search strategy allows the user to select specific terms from an assortment of characteristics that have been used to classify books in the databases. It will be used to demonstrate additional aspects of interaction with the BookHouse system.

As a result of clicking on the person-activity icon that corresponds to the analytical search (i.e., the person seated at the desk) the user adopts the perspective of that person – and sees a number of icons arrayed on and around the desk. Each of the icons contain graphical representations that were designed to symbolically represent a particular dimension that was used to classify the books in the database. For example, the clock on the wall symbolizes the chronological date of a novel. The act of pointing and clicking on an icon is an indication that the user wishes to select an attribute of that dimension for use in searching the database. Upon activation of an icon the user is presented with a graphical depiction of an opened "virtual" book whose pages display a list of keywords. The user can navigate to additional lists by either clicking on one of the "alphabet" icons (to jump to a distant list with terms organized according to the corresponding letter of the alpahbet) or by clicking on one of several arrows (e.g., to go to the next or the previous list).

The user specifies a key word for the search by pointing and clicking at an item on the virtual page. This action has several consequences. First, the key word appears in a text box at the bottom of the screen, indicating that it is now a term that will be used to execute the search. Second, additional icons appear at the top right of the screen, indicating that other actions may now be taken (in this case these are actions that can be used to modify the search). Third, the search is executed. Fourth, the results of the search are portrayed as an icon (i.e., a set of books) that appear in the lower right hand corner, along with a number that indicates how many books met the search criterion.

The user has several options at this point. The results of the search can be examined by clicking on that icon. A verbal description of the first book appears as an entry on the pages of a virtual book. This description is organized according to the dimensions of classification that are used in the system; the description provides a detailed explanation of its contents (and in the process, just how this particular novel fits into the general classification scheme). Additional books meeting the search criteria can be examined by pointing and clicking on the arrows. The information about a book can be printed out (by clicking on the printer icon) or placed aside for future reference (by clicking on icon with the hand holding a book).

Figure 8. Another "scene" from the Bookhouse. The objects represent tools for specifying dimensions in the AMP scheme form categorizing fiction. These tools can be used to "analytically" narrow the search to identify potentially interesting books.

Alternatively, a second primary option is to refine the criteria used in the search. Imagine that the user wanted to narrow down the results of the search by including a keyword from another classification dimension. To do so the user would point and click at the iconic representation of the analytical room located at the top and middle section of the interface. This returns the user to the analytical room; he/she would then follow the steps outlined above to specify a keyword in the second dimension. This keyword is then added to the list at the bottom of the page and the system automatically executes a second search of the database using these two terms combined with a boolean "and."

One other fundamental aspect of the interface needs to be mentioned. A primary consideration in the design of the interface was to allow users to switch easily and seamlessly between databases and search strategies. There are a number of interface conventions that allow this to occur. A primary convention is the strip of navigational icons that appears across the top of the interface as the user navigates through the virtual library. For example, the user can easily choose an alternative database to search by clicking on the "database" icon in the navigation strip (second icon to the right) and then making the approiate choice. Similarly, the user can easily choose an alternative search strategy by clicking on the "strategy" icon in the navigation strip (the third icon to the right) and selecting another person-activity icon. In some cases the possibility for alternative search strategies is even more direct. For example, when the results of a search are examined (i.e., when a book description is on the screen) the user can immediately switch to a search by analogy for the displayed book if they click on the icon with the approximation symbol in the upper right.

The spatial metaphor of navigating through rooms and selecting icons was chosen to make the syntax of the interaction as intuitive as possible. However, it is important to keep in mind that the "ecological" aspect of this interface has more to do with the scheme for classifying books, than with the spatial metaphor. It is in the classification scheme that the question of meaning is directly addressed.

Example 2: The Multidimensional Scaling Interface

MDS-Interactive is an interface designed to help people to explore digital product catalogs. Originally its design was developed to support designers querying a visual database of existing products in the early phases of a design problem. Designers often review precedent designs for form, aesthetics, style, and solution elements. But later it was realized that there are many situations in which people have to choose products in a domain with which they are not familiar. The solution for designers was found to hold well for other problems such as searching for a new home from the haphazard collection of houses on the market, finding a gift, selecting a holiday from among many options, or finding a TV show to watch.

MDS-Interactive was developed to counter some of the problems that designers (and consumers) have with regular database systems (Stappers & Pasman, 1999). (1) The form of the interaction is mostly verbal, through discrete symbolic commands (e.g., search strings). (2) The user must be familiar with the jargon used in describing the content of the database, and must be able to phrase the goal of the search in terms of this jargon. (3) The hierarchic form of the database often imposes an order in which different parts of the question must be addressed. (4) It is easy to lose track of your search steps and your decision process; for example, when using a web search engine, many people find it difficult to retrace their steps or even to remember the decisions they took on an item four steps back. To counter these problems, MDS-Interactive uses a mainly visual interaction style, by letting the user interact with a small set of samples from the database, and presenting the samples in the form of an interactive map which presents similarity of samples by proximity of the sample's images.

In this section we illustrate MDS-Interactive by stepping through two small scenarios. Then we discuss the elements in the interface, and the technical underpinnings of the method (the left and right halves of Figure 9, respectively), indicating how and where it fits in the scheme of the theoretical sections above.

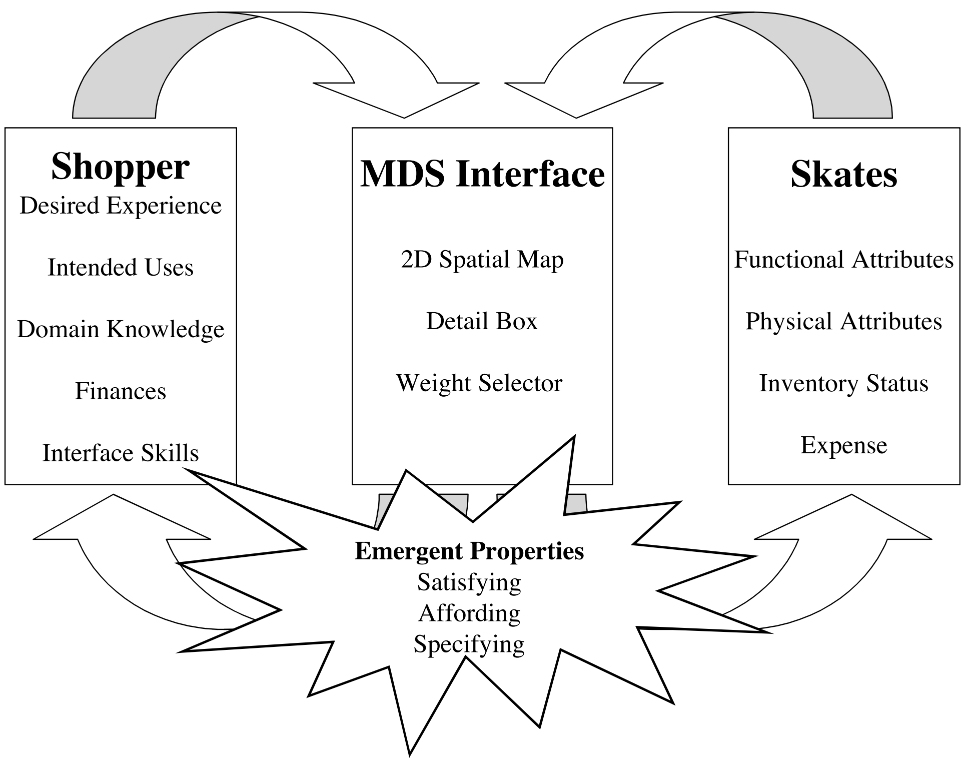

Figure 9. The challenge for the MDS interface was to help users with a vague notion about what they want (in terms of both what's available and implications of difference relative to functional goals) find a product that they will like.

Scenario 1: 'finding a gift for Joey.' Mark wants to buy a gift for his nephew Joey's 10th birthday. Although he knows the boy, he doesn't have a precisely described search question. He turns to an Internet shopping site using MDS-Interactive, and it shows him a small set of gift samples, among which a teddy bear, a football, and a remote-controlled car. As Joey likes to play outside, he clicks near the football and receives some similar sample products: an inline roller skate, a skipping rope, and a fishing rod. Mark likes the skates most, drags the other samples off-screen, and calls up some more roller skates by clicking near the skates. The MDS-Interactive screen now depicts a field of sample skates, which cluster into groups. Some groups appear to be visually similar (e.g., the off-road skates have only two or three wheels instead of the regular four and the speed skates have five). The visual groupings and the appearance of the skates convey some of these aspects directly, or invite Mark to find an explanation of the grouping by inspecting the properties of each sample: as the mouse pointer rolls over a sample, the product info window shows a high-resolution picture of that sample, and lists a number of the key attributes about the sample. By comparing the elements in a group, Mark finds out that all the two-wheeled skates are used for off-road skating.

Mark continues searching the database, removing uninteresting samples (by dragging them out of the circle), and calling up new samples in interesting parts of the sample map (by clicking at interesting positions, near desirable samples and away from undesirable ones). Mark narrows down the sample collection to find 'a good gift for Joey'. Along the way his originally vague goal gets more and more clarified, and Mark learns about the collection of skates in the database, and their properties. Finally, he settles on a 'best gift sample'. In the process he will have considered properties as diverse as 'application area', 'visual appearance', brand, quality and price. He will probably have shifted search strategy a few times, at some times relying on the pictures of the samples, or the grouping of the samples, at other times comparing samples on their properties.

Scenario 2: 'Find the grind.' Joey has been skating for a year or so, knows a bit about skating, and wants to get a new pair, but one that is better for cool stunt work. In MDS-Interactive's initial display, he quickly weeds out the skates he doesn't want, e.g. speed skates and off-road skates, and focuses on regular, compact four-wheeled skates. He is especially keen on questions of how the skates are used, so he sets the weight selector toward 'application'. As a result, the samples in the display arrange themselves on the basis of the properties in the database that have to do with 'application': properties of price and brand ('quality') and the type of fixture and appearance ('design') are ignored. The arrangement now reflects different clusters. The property list helps him to identify the groups as 'middle of the road', 'all-round', and 'grind' skates.

He finds out (by asking the database or from sources outside) that 'grind' is the jargon for doing stunts. After eliminating all non-grind skates, Joey resets the weight selector to bring back aspects of quality and design. Now he compares the collection of grind skates. By studying the remaining clusters, their property lists, he will find also that there are considerable price differences (which often concur with brands), and that there are differences in the 'ball bearings' within this group. Prompted by this difference, he calls up information about ball bearings and finds out that these differ in hardness and price, and that grinders care a lot about the quality differences in these.

In each of the two scenarios we have a competent user conducting a search by negotiating his way through a set of samples. During the search, new samples and the way in which they are arranged implicitly prompt the user to make rough decisions in the visual display, to look deeper into the samples' properties, or to change the relative weights of the different properties. All these actions occurred 'out in the open', by direct manipulation of the samples, the space between the samples, and item property window and the property weight selector.

As the user goes along, he or she not only learns about the skates in the database, but also learns to relate the objective of the search to the jargon terms that are used in the database. Typically the users learn about the jargon only when it affects their decisions. For example, if all the skates that the user is considering have the same quality ball bearings, there is no need for him or her to learn the ins and outs about ball bearings. This type of search occurs in many product choice decisions: when someone wants to buy a house, he often cannot phrase the question in terms as "it must have 4 rooms, one of 5' x 12', and a kitchen with two windows", even though these are relevant terms for deciding if a certain house fits. The searcher for a house is often looking for a best fit for a complex set of considerations, not for an aggregate of *optimal properties*, but for an *existing* sample that fits better than the others.

Domain. As with the Bookhouse, the graphical interface is the most salient surface feature of the MDS design. However, as with the Bookhouse the most important dimension, from the perspective of CSE is the domain structure that can be accessed through this interface. What differences make a difference in terms of satisfying, affording, and specifying? At the top level of the search we have the product samples (which evolve from ‘gift’ to ‘outdoor products’ to ‘skates’ in Scenario 1). This top level is found in the field of samples. Below that lies the level of aggregate skate properties ‘quality’, ‘design’ and ‘application’, which have direct implications for whether a particular skate will satisfy the goals of the users (in terms of use of the skates and quality of the experience). Next there are the properties that identify the skate relative to other skates in the inventory (e.g., price and brand). And finally, there are the physical materials, such as the quality of ball bearings.

At the level of skates, much of the structure comes from physical inventory of skates and the intuitions of experts about what dimensions of this inventory make a difference in terms of how skates fit various uses or applications. The roller skate prototype contained 70 skates, described by 20 properties, organized in 3 groups, which are shown in the weight selector. The layout algorithms used all 20 properties, but only a few were readily available in the detail window. The properties were selected and their relative weights were set by a domain expert, so that the layout patterns were based on similarities that could be expected to be “meaningful” to the users of the interface.

At the level of an interesting gift for Joey – the structure comes from the user’s experience with Joey. Building this into the interface structure would require a theory of what makes things interesting and perhaps a theory of “Joey.” Without such a theory – it is still possible to “let” these constraints enter the process by enabling the searcher to explore freely (in other words, to let the searcher browse with minimal interference). By minimizing constraints associated with the interaction – we maximize the opportunity for the system to organize around constraints that the human brings to the system (e.g., such as a concept of what Joey might like).

Agency. As with the Bookhouse the target user population is a very heterogeneous group, both in terms of what they might be looking for and in terms of the skills they have for navigating through the Web. Again, a spatial metaphor is seen as a good solution because it taps into the general experiences of moving and manipulating objects that many people will bring to the interaction.

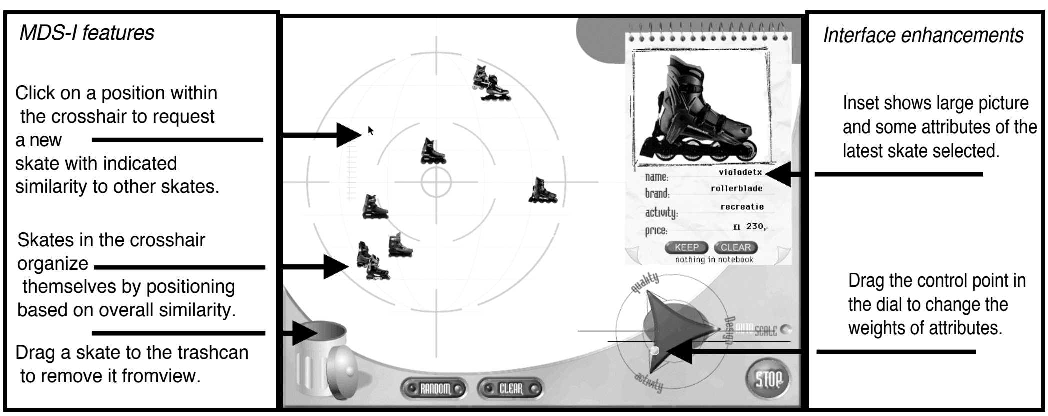

Figure 10. The interface for MDS-Interactive which allows search through a database of skates to discover one that satisfies your personal needs.

Ecological interface. The interface shown in Figure 10 consists of three components:

(1) A field of samples whose layout is continually adjusted to reflect their overall similarities. When new samples enter the field, old samples are removed, or the user tries to drag a sample to another location, the field adjusts and the samples settle to a new layout. The mathematical technique used to determine the optimal layout from the properties of the samples is called Multi-dimensional Scaling (MDS), a method that has been used in explorative statistics since the 1960s (see, e.g., Kruskal & Wish, 1978; Borg & Groenen, 1997). What is new in the interface is first its dynamic nature (new arrangements are reached by animations showing the motions of all samples), and foremost its interactive nature (Stappers, Pasman, & Groenen, 2000). When the user clicks between the samples, a database query is performed to find the sample in the database which best matches the place of the click, i.e., the sample that is very similar to the

samples that are close in the field, and at the same time dissimilar to the samples that are on a larger distance. Through this, the field of samples becomes a field of meaning, where each position in the field expresses a query or a result.

(2) a detail window listing a set of property values of the currently selected sample.

(3) a weight selector which lets the user emphasize some properties and de-emphasize others properties. In this way the display (and the queries) can be directed to a specific aspect.

This interface allows representation of meaning at many levels. At the top level of the search there are the product samples (which evolve from ‘gift’ to ‘outdoor products’ to ‘skates’ in scenario 1). This top level is found in the field of samples. Below that lies the level of aggregate skate properties ‘quality’, ‘design’ and ‘application’, which can be manipulated with the triangular weight selector. Finally, a selection of the properties that are expected to have a direct meaning for the user, such as price and brand, are shown in the detail inset. But most low-level properties, such as the materials of the ball bearings, are not normally on the display. Yet, as described in the second scenario – the hope is that attention will be drawn to these properties as they become meaningful in the course of the interaction.

The MDS-Interactive interface has been user-tested for a small range of application areas, product types, property types, ranging from whiskies (classified on subjective taste judgments), through consumer products such as roller skates (integrating technical properties and experiential judgments), to information services, such as a digital TV-guide (classifying programs based on formal classifications and properties such as time and length of broadcast and type of audience appeal), and expert information systems (such as marketing evaluations of company portfolios). Findings from user-testing indicate that the simplicity of the interface, and the fact that all of the interactions are visual and in the open, enable people to explore collections starting out with vague, taste-based, or expert-based questions. Moreover, people could quickly learn how to search by watching one other person use it, and people could use it in collaboration (whereas most search or browse systems provide very little support for collaborative use).

Summary and Conclusions

> *One can think of problem solving situations as composed of three basic elements; the world to be acted on, the agent who acts on the world, and the representation of the world utilized by the problem-solving agent. . . . There is an intimate relation between these three elements so that cognitive description is at once a psychological description and a domain description. . . . a complete description must specify the relationships between all three sets of factors.* Woods & Hollnagel, 1987, p. 258)

Among those concerned about the “human factor” in complex systems, you will find almost unanimous agreement with this statement of Woods and Hollnagel. Thus, there is not much controversy about the “span” of the problem. However, there are many disagreements about how meaning fits into this picture. What is the source of meaning? What is the nature of the nature of the underlying dynamic (e.g., the nature of causality)? What is the appropriate conceptual parsing for making sense of this dynamic? What are the implications for design? Our goal in this chapter was to provide an “ecological context” for framing these questions. For us some distinctive features of an ecological approach are:

1) The focus on the ecology (or in James terms “experience”) as the source for meaning. This approach acknowledges the contributions of both animal and environment to meaning – but gives priority to neither.

2) A “circular,” as opposed to a “linear” view of causality. This circular dynamic, depicted in Figure2, reflects a coupling between perception and action and between control and exploration. Within the circular dynamic neither perception or action, or control or exploration takes precedence over the other. Each occurs by virtue of the other. When these processes are coupled in a symbiotic fashion the result is normally an improving “fit” between organism and environment.

3) In parsing this dynamic, the attributes of “specifying,” “affording,” and “satisfying” emerge as the objects of study. These are seen as critical dimensions of the “fit” of organism and environment that cognitive science should be exploring – both as a means to a richer cognitive science and as a bridge

between basic questions about human performance and practical questions about designing to support meaning processing.

4) The foundation for design is first to ask the question, what is meaningful. What differences make a difference in terms of the functional goals of the work domain? The next question is to explore ways that these differences can be made more explicit to the humans in the system – either through the design of ecological interfaces, training, or both. We tried to provide a logic case for this in the discussion of CSE and tried to illustrate the design problem with the Bookhouse and the MDS-Interactive Interface.

In conclusion, we hope that this chapter will be appreciated by others, who are struggling to understand the nature of cognition and the implications for design, as an interesting alternative to more conventional approaches – not an answer – but an interesting perspective. We also hope that this chapter will dissuade those who are looking for simple answers to complex problems from identifying their work with an ecological approach. The ecological framework is an attempt to scale-up theory to address the complexity of work environments. Not an attempt to reduce the problems of work into a convenient theoretical paradigm. It is an attempt to make science and design "phenomenon" or "problem centered." Above all, it is a search for meaning!

References

Borg, I. & Groenen, P.J.F. (1997) Modern Multidimensional Scaling: Theory and Applications, Berlin: Springer.

Dreyfus, H.L. (1994). *What computers still can't do: A critique of artificial reason.* Cambridge, MA: MIT Press.

Fischer, G. (2000). "Design, Learning, Collaboration and New Media--A Co-Evolutionary HCI Perspective." In C. Paris, N. Ozkan, S. Howard, and S. Lu (Eds.), *Proceedings of OZCHI 2000*, Sydney, Australia, December, pp 282-289.

Flach, J.M., Jacques, P.F., Patrick, D.L., Amelink, M., van Paassen, M.M., & Mulder, M. (In press). A search for meaning: A case study of the approach-to-landing.

Gibson, J.J. (1972/1982). The affordances of the environment. In E. Reed & R. Jones (Eds.). *Reasons for realism: Selected essays of James J. Gibson*. Hillsdale, NJ: Erlbaum.

Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.

James, W. (1909). *The meaning of truth: A sequel to Pragmatism*. New York: Longmans, Green.

Kruskal, J.B., & Wish, M. (1978) Multidimensional Scaling. Beverly Hills,, CA: Sage.

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

Pejtersen, A.M. (1979). Investigation of search strategies in fiction based on an analysis of 134 user-librarian conversations. In T. Henriksen (Ed.) *IRFIS 3 Conference Proceedings* (pp. 107-132). Oslo Norway: Staten Biblioteks-och Informations Hoegskole.

Pejtersen, A.M. (1984). Design of a computer-aided user-system dialogue based on an analysis of users' search behavior. *Social Science Information Studies*, 4, 167-183.

Pejtersen, A.M. (1988). Search strategies and database design for information retrieval from libraries. In L.P. Goodstein, H.B. Andersen, & S.E. Olsen (Eds.). *Tasks, errors, and mental models: A festschrift to celebrate the 60th birthday of Professor Jens Rassmussen.* (pp. 171-190). London: Taylor & Francis.

Pejtersen, A.M. (1992). The Bookhouse: An icon based database system for fiction retrieval in public libraries. In B. Cronin (Ed.) *The marketing of library and information services 2* (pp. 572-591). London: ASLIB.

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

Rasmussen, J., Pejtersen, A.M., & Goodstein, L.P. (1994). *Cognitive Systems Engineering*. New York: Wiley.

Simon, H. (1981). *The sciences of the artificial*. (2nd ed.) Cambridge, MA: MIT Press.

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

Schön, D. A. (1983). *The reflective practitioner: How professionals think in action.* New York: Basic Books.

Stappers, P.J., & Pasman, G. (1999) Exploring a database through interactive visualised similarity scaling. In M.W. Altom & M.G. Williams (Eds.). *Human Factors in Computer Systems. CHI99 Extended Abstracts.* (pp. 184-185) May, Pittsburgh, PA.

Stappers, P.J., Pasman, G., & Groenen, P.J.F. (2000) Exploring databases for taste or inspiration with interactive multi-dimensional scaling. In *Ergonomics for the New Millenium. Proceedings of the XIVth tirennial confress of the International Ergonomics Association and 44th annual meeting of the Human Factors and Ergonomics Society*, San Diego, CA, July 29-August 4, 2000, Pages 3-575 – 3-578.

Vicente, K.J. (1999). *Cognitive work analysis.* Mahwah, NJ: Erlbaum.

von Uexküll, J.(1957). A stroll through the worlds of animal and man. In C.H. Schiller (Ed.), *Instinctive behavior* (pp. 5–80). New York: International Universities Press.

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

Woods, D.D. & Hollnagel, E. (1987). Mapping cognitive demands in complex problem-solving worlds. *International Journal of Man-Machine Studeies*, 26, 257-275.