AMES GRANT
IN-54-CR
69290
P 75

Acquisition and Production of Skilled Behavior in Dynamic Decision-Making Tasks:

Semiannual Status Report

for

NASA Research Grant NAG2-656

January 31 1992

Principal Investigator:

Alex Kirlik
School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0205

(NASA-CR-189846) ACQUISITION AND PRODUCTION OF SKILLED BEHAVIOR IN DYNAMIC DECISION-MAKING TASKS Semiannual Status Report (Georgia Inst. of Tech.) 75 p
CSCL 05H G3/54
N92-17132
Unclas
0069290

Comparison of Display Enhancement with Intelligent Decision-Aiding

Proposed Research

and

Preliminary Analyses

by
Alex Kirlik
Wendy J. Markert
Merrick Kossack

Center for Human-Machine Systems Research
School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332
(404) 894-2373

February 3, 1992

ABSTRACT

Currently, two main approaches exist for improving the human-machine interface component of a system in order to improve overall system performance—display enhancement and intelligent decision-aiding. Each of these two approaches has its own set of advantages and disadvantages, as well as introduce its own set of additional performance problems. These characteristics should help identify which types of problem situations and domains are better aided by which type of strategy. This report first describes the characteristic issues of these two decision-aiding strategies. Then differences in expert and novice decision-making are described in order to help determine whether a particular strategy may be better for a particular type of user. Finally, research is outlined to compare and contrast the two technologies, as well as to examine the interaction effects introduced by the different skill levels and the different methods for training operators.

Introduction

Currently, two main approaches exist for improving the human-machine interface component of a system in order to improve overall system performance—display enhancement and intelligent decision-aiding. These two approaches stem from the two main ways in which to aid human performance. People can be aided in what they perceive (by making important information more easily identified) as is the chief concern of display enhancement, or they can be aided in what they do with the information they perceive (making it easier to perform operations, etc.) which is a goal of intelligent decision-aiding.

These two technologies have their own sets of advantages and disadvantages, as well as introduce their own sets of error modes (or new performance problems). Each approach has characteristics which should help determine which problem situations are better handled by one than the other, as well as what type of user (in terms of amount of experience) is better aided by which type of human-machine interaction.

This research intends to compare and contrast these two technologies—display enhancement and intelligent decision-aiding—to determine which types of human-machine systems are better improved by each. Specifically, the domain examined here is a complex, dynamic decision-making task, but it is hoped that a thorough study will lead to conclusions which can be made about other domains as well. A look will also be taken to determine if different skill levels of users are aided in different ways by these two approaches. Also, the possible error modes introduced by each type of aiding will be examined.

This report is divided into five main sections. The first section will examine the characteristic issues of intelligent decision aids, as well as determine the possible error modes introduced by this type of aiding. The next section will discuss the issues surrounding the technology of display enhancements, as well as determine the possible error modes introduced by this technology. The third section will identify some of the differences in expert and novice decision-making which may help identify how different types of aiding systems may aid different skill levels in different ways. Then research will be outlined in which to compare and contrast the technologies of intelligent decision-aiding with display enhancements, as well as the effects found when using both with different skill levels of operators. Finally, the expected contributions of this research will be described.

Characteristic Issues of Intelligent Decision-Aiding

Intelligent decision-aiding is an area which has been the focus of much research in the hopes of discovering ways to improve the performance of human-machine systems. Woods [1986], Woods and Roth [1988], Norman [1988], and Hopkin [1988] are just a few of the many researchers who have tried to determine what makes a good intelligent decision aid.

Among the issues which need to be addressed in the design of such a system is the role of the decision aid in the system. Woods and Roth [1988] discuss the role of the aid as an instrument (as opposed to as a prosthesis which is meant to replace people or to fill a human deficiency). As an instrument the decision aid is to be used as a reference or extra source of information for the human problem solver. This perspective leaves the human in control and performing more than mere supervisory tasks. People are not particularly good at monitoring, and therefore, keeping the controller active is an advantage. The system is also more flexible when the aid is used as a tool and not as a replacement, in that special cases can be more easily handled.

A decision aid can merely identify the existence of a problem or it can also give advice to help in solving the problem. Woods and Roth [1988] define good advice as more than recommending a solution. Advice needs to be given in the situations which call for it, and only when needed. If failures of attention are a problem in performing a task, an aid which focuses the operator's attention on the relevant information is an appropriate goal for the designer [Woods 1986].

The issues of usability and understandability need to be addressed. Information needs to be presented in clear formats that are easy for people to use and understand. After all, what good is providing the user with a decision aid if the user can not use it properly or does not understand what the aid is telling him? To guard against this problem, Norman [1988] advocates doing "user-centered design." He suggests the following principles of good design:

- Make things visible: The user should be able to tell what is going on by merely looking at the interface.
- Provide the user with a good conceptual model: The designer should provide the user with a mental model that is consistent and coherent.
- Provide the user with good mappings: The user should be able to determine the relationships between actions and their results.
- Provide the user with feedback: The user should receive full and continuous feedback about the results of all actions.

Another key issue is that of user acceptability. What good is a decision aid if the user does not want to use it or does not have faith in it? If the user is expected to override the computer recommendations, he should have the actual authority to do so. If the aid is used continually, care needs to be taken so that boredom is not aggravated in users when the work load is light, by reducing their workload even more [Hopkin 1988]. Often dangerous accidents and failures of attention occur when the user is less active, rather than under high workload situations.

The type of information which is revealed to the operator is another consideration of this type of aiding. If the operator receives just a warning indicating a problem, but is not told what the problem is, this is not very useful and could be found to be more of a nuisance than an aid. Also important is how the operator is alerted to a warning. For example, if he hears an annoying buzz or beep every time there is a possible conflict, and if this were to happen constantly, the buzz or beep could prove more distracting and a nuisance, and may in fact no longer serve its function if the user decides to tune it out (or even to disable it) [Norman 1988].

Also important to consider is whether the aid is working at the correct level of abstraction or whether it is supporting the correct mode of problem solving. Vicente and Rasmussen [1990a, 1990b] propose a method for interface design they call ecological interface design (EID). EID is based upon Rasmussen's skills, rules, and knowledge (SRK) framework for human performance and on his means-ends abstraction hierarchy which illustrates the functional properties of a system [Rasmussen 1986]. The SRK framework proposed that information can be detected in three ways—as signs, signals, and symbols. Information is perceived as signals when the operator detects the time-space behavior of the data. Signs are interpreted when the perceptual characteristics of the data are detected. Finally, symbols are perceived which represent concepts and have meaning. SRK claims that the way in which information is detected is related to the way it is processed [Vicente 1988]. Therefore, the three different forms if information—signals, signs, and symbols—refer to three different processing modes—skill-based, rule-based, and knowledge-based behavior. "The basic implication is that one should design interfaces in such a way as not to force cognitive control to a higher level than the demands of the task require, while at the same time providing the appropriate support for all three levels [Vicente and Rasmussen 1990a]." By providing a sentential aiding system, skill-based performance is not supported. Signs and signals are not conveyed, only symbols—the words.

The abstraction hierarchy is a tool used to describe the functional properties of a system. "Such a hierarchy describes bottom-up what components and functions can be used for, how they may serve higher level purposes, and, top-down, how purposes can be implemented by functions and components [Rasmussen 1986]." The levels in the abstraction hierarchy include physical form physical function, generalized function, abstract function, and system purpose. The hierarchy provides a way to structure the properties to be represented in the interface, while the constraints revealed within the hierarchy enable the operator to focus his attention on the most appropriate system component by crossing through various levels of the hierarchy. In terms of the level of abstraction, different levels can be supported by the intelligent decision aid, depending on exactly which information is presented to the user. However, in order to support all levels in the abstraction hierarchy, it would probably be necessary to supply the operator with so much verbal information that it would require too much time for him to use it properly and it would require more space than desired to implement.

Another concern is what happens if the aid is inoperational and the operator needs to function on his own. Will he lose the skill he needs to detect possible problems on his own after using an aid for a while [Woods 1986]? If the user is depending on the aid to alert him to dangers, will he stop using particular cues which he would need to perform this replaced function, but which would also help him in determining a resolution to the problem?

Therefore, we see that the following possible error modes may be introduced into a system through the use of intelligent decision-aiding:

- increased boredom of the operators leading to worsened performance;
- lack of trust in the aid;
- lack of responsibility taken by the user for override of the decision aid;
- skill reduction in the operator;
- failure of the operator to attend to important situational cues;
- lack of user acceptance of the aid;
- inability of the operator to identify problem resolutions (when they are not provided by the aid);
- worsened system performance due to the aid being a nuisance.

The next section considers the issues surrounding enhanced displays. As will be seen, many of the issues are similar to those discussed in this section. However, due to the inherent nature of these two technologies, they each have a set of issues which characterize them with respect to system domain and implementation.

Characteristic Issues of Display Enhancements

Display enhancements have been more recently proposed as an alternative to intelligent decision aiding for certain human-machine systems. Hammond, Hamm, Grassia, and Pearson [1987] and Vicente and Rasmussen [1990a, 1990b] have looked at the effects which interface design can have on modes of processing and at how interface design can be improved so as to take better advantage of natural human abilities, two of the issues involved in determining how enhanced displays can work as a tool for decision aiding.

Because of the inherent nature of enhanced displays the role of this technology in the system is that of an instrument. Again, the human operator is left in charge while the enhanced displays are designed to aid him in focusing on the correct information when it is needed. However, this is the area where the designers need to be the most careful. The question is whether, when actually implemented, the enhancements actually help the user to find, integrate, or interpret the right data at the right time. If implemented incorrectly, the enhancements can create confusion and become a hindrance. Larkin and Simon [1987] in studying pictorial representations, have found that although the following characteristics are not sufficient for a diagram, or in this case an interface display, to be useful, they are necessary for the construction of a useful pictorial representation: information to be used together should be grouped together in order to reduce the search required to find the necessary elements; location should be used to group information about particular elements to eliminate the need to match symbolic labels; and perceptual inferences should be supported. Along with these guidelines, the suggestions made by Norman [1988] for effective interface design should also be followed. Related to this issue is again the issues of usability and understandability, as discussed in the previous section.

Another concern is whether the information displayed through the enhancements actually does aid the operator in making a decision. When a conflict or potentially dangerous situation arises, do the displays merely inform the user of this possibility, or do they also help him in deciding how to resolve the problem situation? Since Woods and Roth [1988] have defined good advice as that which is given in situations which call for it, and *only when needed,* will enhanced displays create more confusion by conveying extra information, or "advice," continually?

Woods [1986] also recommends for goal-directed knowledge representation, "possible actions are organized around the goals that they can effect, and data about pre-conditions, post-conditions (effects), constraints (side effects or inter-goal couplings) and alternative means are captured." In other words, entire contexts are considered. However, in a complex, dynamic system, the current "context" is always changing. It is possible that the display enhanced interface will not be able to isolate contexts for each task, and in doing so does not focus the operator's attention on the important information. However, Woods [1986] also claims that "if available data are organized and displayed so that the user can directly see the state of task-meaningful objects [as is the case in general for display enhanced interfaces], then natural mechanisms for focusing in on the relevant data for the current context will be more effective."

In terms of user acceptability, because of the inherent nature of enhanced displays they are different from intelligent decision aids in how the user views them as part of the system. To the user, the enhanced displays are merely a part of the user interface of the system he is working with. The display is not a computer trying to tell the person what to do (although indirectly they are). In this sense, user acceptability is viewed differently in considering enhanced displays then when considering traditional decision aids. If the user is content to work with the system interface as presented to him, the designer needs not consider what would happen if the user decides to ignore direct suggestions for actions or warning of potential crises—instead the user sees himself as determining these for himself.

Similar to intelligent decision-aiding is the question of what happens if the aid (in this case the system with the enhanced displays) is inoperational and the operator needs to function on his own. Will the user lose the skill he needs to detect possible problems on his own after using the aid for a while? Or will the use of the aid have effectively trained him to know where to look on the system display by himself? This issue is important not only in the case of an emergency or unexpected situation, but it brings up many points related to the training of operators, as well.

Finally, another issue to consider is whether the display enhancements are at the wrong level of abstraction or support the wrong mode of problem solving. The abstraction hierarchy and the modes of problem-solving associated with the SRK framework were described in the previous section. Enhanced displays have the inherent capability of better supporting these three processing modes than the decision-aiding system alternative when they are designed properly. Also, in terms of levels of abstraction, the enhanced display version of a decision aid better supports such a structured view if the system properties to be represented in the system interface by providing support for answering the questions, WHY (going up in the hierarchy) and HOW (going down). The interface for the decision aiding system does not directly support these questions or this structure.

Therefore, from this discussion we can see that the following possible error modes may be introduced into a system by the introduction of enhanced displays:

- skill reduction in the operator;
- failure of the operator to attend to important situational cues;
- inability of the operator to identify problem resolutions;
- worsened system performance due to the aid creating more confusion and being a hindrance.

Now that we have examined the issues characteristic of each type of decision-aiding technology, it is important to investigate the differences in decision-making processes found in varying levels of expertise in operators. These differences may help to identify what is needed of a decision aid intended to help a particular type of operator. They may also lend advice when determining how to best aid the training of operators.

Differences in Expert and Novice Decision-Making

Psychologists have been interested in studying the distinctions between expert and novice behavior in many different task domains. For example, Chi, et. al. [1981] have studied the differences in solving physics problems, Staszewski [1988] has looked at expertise in mental calculation (specifically, "lightning mental calculators"), and Soloway, et. al. [1988] have examined how expert and novices write (and read) computer programs. Tasks like these are concerned exclusively with the acquisition of cognitive skills. This work contrasts with that which focuses on skill acquisition of perceptual skills. It is the latter which will be examined in this section because the the research proposed here will involve human-environment interaction.

Investigating how people interact with their environment, Dreyfus, Dreyfus, and Athanasiou [1986] have witnessed five common stages of skill acquisition in the progression from novice to expert. Their research has dealt with such areas as piloting airplanes, playing chess, driving automobiles, and learning a second language. The five stages they found are novice, advanced beginner, competent, proficient, and expert.

The authors theorize that this progression from novice to expert is characterized by the following behavior. During the novice stage context-free elements and rules are learned. These elements and rules are considered context-free because they are clearly defined and easily recognized independent of the specific situation in which they occur of apply. Eventually, with practice and increased experience, novices learn to recognize "situational" elements which are context-sensitive and not objectively definable. At this point they become advanced beginners. More experience leads to the ability to adopt a plan to organize a specific situation. In this way competent behavior allows someone to improve his performance and make decisions in a hierarchical manner. Proficient behavior is characterized by the ability to "intuitively" make decisions be relating present situations to previously experienced situations. In this way the person can use expectations and previously used plans to solve new problems. Finally, an expert makes decisions and solves problems easily and without effort. He performs naturally and does "what normally works."

This model has not yet been formalized (computationally) and has not yet been rigorously evaluated. However, it does provide a way of looking at the progression from novice to expert which considers the environmental factors and which appears to occur across a wide variety of domains.

Focusing on the final stage of acquisition, Klein [1989] has proposed a model to explain expert decision-making which consists of four main steps: 1. recognizing cases as typical; 2. understanding the situation; 3. evaluating alternatives; and 4. progressive deepening. Decision makers rely on their previous experience to recognize and classify situations as typical. Understanding a situational is comprised of recognizing four types of information—plausible goals, critical cues and causal factors, expectancies, and typical actions. Once classified, the decision maker can recall the typical way of handling that type of situation. He would use available time to evaluate whether or not an option would be appropriate. This might be done using imagery, where the decision makers would imagine implementing the option, in order to determine if anything might go wrong. If problems do arise in this imaging, the plan could be modified or even rejected. If a satisfactory plan if found, it is implemented. If a plan is rejected, a new one is selected and evaluated.

While the cognitive model here consists of familiar situations, typical actions, goals, expectancies, progressive deepening (imagining how an option will be carried out within a specific situational context), evaluation, and selection, the environmental model consists of situational cues, actions, decision points, and outcomes. Therefore, the model is sensitive to the critical cues which the decision-maker has learned to recognize. The expert decision maker is the one who has learned to distinguish which of the available cues are the critical ones. Although this work is supported by studies Klein has done in perceptually rich domains involving fireground commanders, tank platoon leaders, and design engineers, it has not been rigorously tested or completely formalized either.

Lesgold, et. al. [1988] have studied varying levels of performance in another perceptually rich domain, that of X-ray diagnosis. They have found that experts first seem to develop mental representations which in turn direct perception. Experts also appear to evoke an appropriate behavioral schema rather quickly. They know where to look and what cues to look for. In other words, they "have the ability to discriminate between relevant information and 'noise' in a given domain of action, by invoking both precepts and practice ... [Suchman 1987]." In contrast to novices, they were able to distinguish subtle differences and were more flexible in considering other possibilities based on feature detection. Lesgold, et. al. also found that performance was a nonmonotonic function of experience (similar to results found in language learning [Hetherington and Parke 1986]), that is, performance does not increasingly improve as people gain more experience. In between the stages of novice and expert, people reach a stage where their performance degrades slightly. Lesgold, et. al. studied aspects of both cognitive and perceptual learning. They proposed that the development of expertise first comes through a "perceptual tuning" in which the stimulus pattern was classified with the diagnostic decision which had the highest probability of occurring. The result of this perceptual processing would then be used for cognitive processing to resolve ambiguity. This cognitive processing can not evolve until perceptual processing has been tuned and, therefore, they propose that the development of expertise is a shift from purely perceptual decision-making to progressively deeper cognitive decision-making.

Using a different approach, DeGroot [1965] was interested in discovering the differences between players of varying degrees of expertise in the domain of chess. Specifically, DeGroot looked at chess players at the expert and Grand Master levels. Each subject was given the identical chess position and then verbal protocols were taken as he decided what move to make. Results showed that there was essentially no difference in the thought processes between the two groups—search patterns, number of moves considered, etc. The only real difference between the levels was in the quality of the move finally chosen.

Similarly, Chase and Simon [1973] conducted a study in which they replicated some of DeGroot's results and tried to isolate and define the structures (i.e., chunks) into which the information perceived by chess players is organized and stored in memory. They used three levels of players—novice, intermediate, and master. Two tasks were used in this experiment—the perception task and the memory task. The perception task required chess players to reconstruct a chess position while the original remained in view. The memory task required players to reconstruct a position from memory after being exposed to the target board for a short time (5-10 seconds). In the perception task, evidence was found indicating more rapid encoding of information for the more experienced players. In the memory task Chase and Simon found the number of pieces per chunk varied among the skill levels and that the pieces within a chunk seemed to have relationships of defence or attack, to be close together, and to be of the same color and type. Finally, an interesting result found was that when both groups of players (experts and novices) were confronted by random positions (not found in actual games), they both did equally poorly.

In contrast to most of the research described here, very little work has been done in domains involving human-environment interaction which are non-adversarial and deal mostly with skill (as opposed to cognitive processes). An example of this type of research is that done by Deakin and Allard [1991] which investigates the ability of expert and novice figure skaters in recalling elements of a figure skating routine. One important finding was evidence that figure skaters do not simply memorize the sequence of elements for a routine in the same manner that they would memorize a verbal list. Also, "expert skaters seem to have faster access to semantic memory for skating elements than do novices."

Allard has also looked at expertise in sports requiring more open skills (occurring in a moving and changing environment such as volleyball [Allard and Starkes 1980] and basketball [Allard, Graham and Paarsalu 1980]). She has found that there are many similarities in chunking and categorizing performance (thought to show the expert's ability to classify elements according to the significance to the situation) between experts in these types of sport domains and those in more cognitive domains such as chess [Allard and Burnett 1985]. Evidence was also found that expertise in this type of domain can be split into two components—declarative knowledge (for cognitive tasks) and procedural knowledge (for playing the game).

Based upon much of this research (especially that done by Dreyfus, et. al. [1986], Klein [1989], and Lesgold, et. al. [1988]) it appears as if novices, who work more with context-free elements and rules and are not as able to identify subtle differences, are working more on a level which can be best described using rules. If this is the case, an intelligent decision aid may make more sense to them because it appears to be making decisions in a way more similar to the processes that the novices themselves use. In contrast, experts behave more intuitively and are very context dependent. Therefore, enhanced displays seem to operate in a way more consistent with how they view the domain. It will be interesting to determine through the experiment proposed in the next section whether this is indeed the case. Also, it will be interesting to see if the different types of aiding have different effects evident in training novices to become experts.

Proposed Research

The research proposed in this section will compare the two decision-aiding techniques of display enhancement and intelligent decision-aiding in the complex, dynamic domain, Star Cruiser (see Kirlik [1990] for a complete description of this task). In particular, we will study the effects that these two aiding systems have in creating possible errors modes, or performance problems, as well as any differences in performance associated with different skill levels of operators.

Decision Aid Descriptions

Although the actual modified versions of the Star Cruiser task have not yet been created, this section is intended to provide the reader with an idea of what the actual changes may involve.

Two modified versions are needed—an intelligent decision-aiding system (DA) and a display enhanced system (DE).

The DA version will provide a traditional sentential type of advice window which will display to the user a ranked list of recommended actions. The recommendations could be ranked, according to how many points could be scored if the recommendation is followed (with more points being better). Also, if the operator tries to execute an action which is undesirable (i.e., deploy a manned ship to a planet which does not support life, deploy a ship to a planet without the appropriate data or resources, load too many data or resources onto the star cruiser, etc.) the system will respond with a dialogue box pointing out the error to the operator. In contrast, the type of enhancements that might be included in the DE version include displaying an ellipse around the star cruiser designating the regions it can travel to and make if back to its star base to refuel, displaying to the operator only those types of ships that can be deployed in the solar system in which the star cruiser is orbiting, highlighting only those craft containing data or resources that can fit onto the star cruiser craft, and highlighting only those planets within orbit that a chosen craft can be deployed to.

Basis for Comparisons

Performance on the two systems can be compared by examining the points scored, the number of sessions which terminated early (due to running out of fuel, trying to load too many data or resources onto the star cruiser, or crashing into suns), the number of times craft were deployed to planets without resources or data, and the number of times manned craft were deployed to planets without life support (and unmanned craft deployed to life-sustaining planets).

Experimental Design

For the main experiment, that of comparing version DA with version DE, nine main subject groups will be required [see Figure 1]. Additionally, different levels of expertise (novice and expert) will also be examined in regard to these nine main subject groups. Appendix A contains the hypotheses that may be tested in regards to the performance of the nine subject groups.

Anticipated Problem Areas

Extra care will need to be taken to try to ensure that neither of the two systems (DA or DE) introduces more information than the other. Also, in order to extract more general conclusions from the results found in this experiment, the Star Cruiser interface will need to be tested to determine how sensitive it is to different implementations.

system trained on

system tested on

	unaided	DE	DA
unaided	**A**	**D**	**G**
DE	**B**	**E**	**H**
DA	**C**	**F**	**I**

Figure 1: Subject groups for the main experiment

Preliminary Analyses

Before the modified versions of the Star Cruiser task described in the previous section can be created, a model of the Star Cruiser environment needs to be developed, which is independent of the actual information contained in the implementations of the interface. An abstraction hierarchy is a means of doing this, without specifying strategies that should be used by the operator. A preliminary hierarchy is shown in Figure 2 (Note: the second level is still in development). Figure 3 shows the further decomposition of the **collection** function.

Also necessary before developing the decision-aiding versions of Star Cruiser, is the identification of the different types of actions available within the Star Cruiser task with which the user may need help in deciding what to do. A preliminary listing of these actions is found in Appendix B. After these actions have been identified, those actions which can be aided both perceptually (through enhanced displays) and cognitively (through sentential advice) need to be discovered, as well as the means for implementing the aiding devices for these actions.

Finally, the chosen implementations for the decision-aiding techniques will be coded in order to create the two modified Star Cruiser versions to be used in the comparison study described previously.

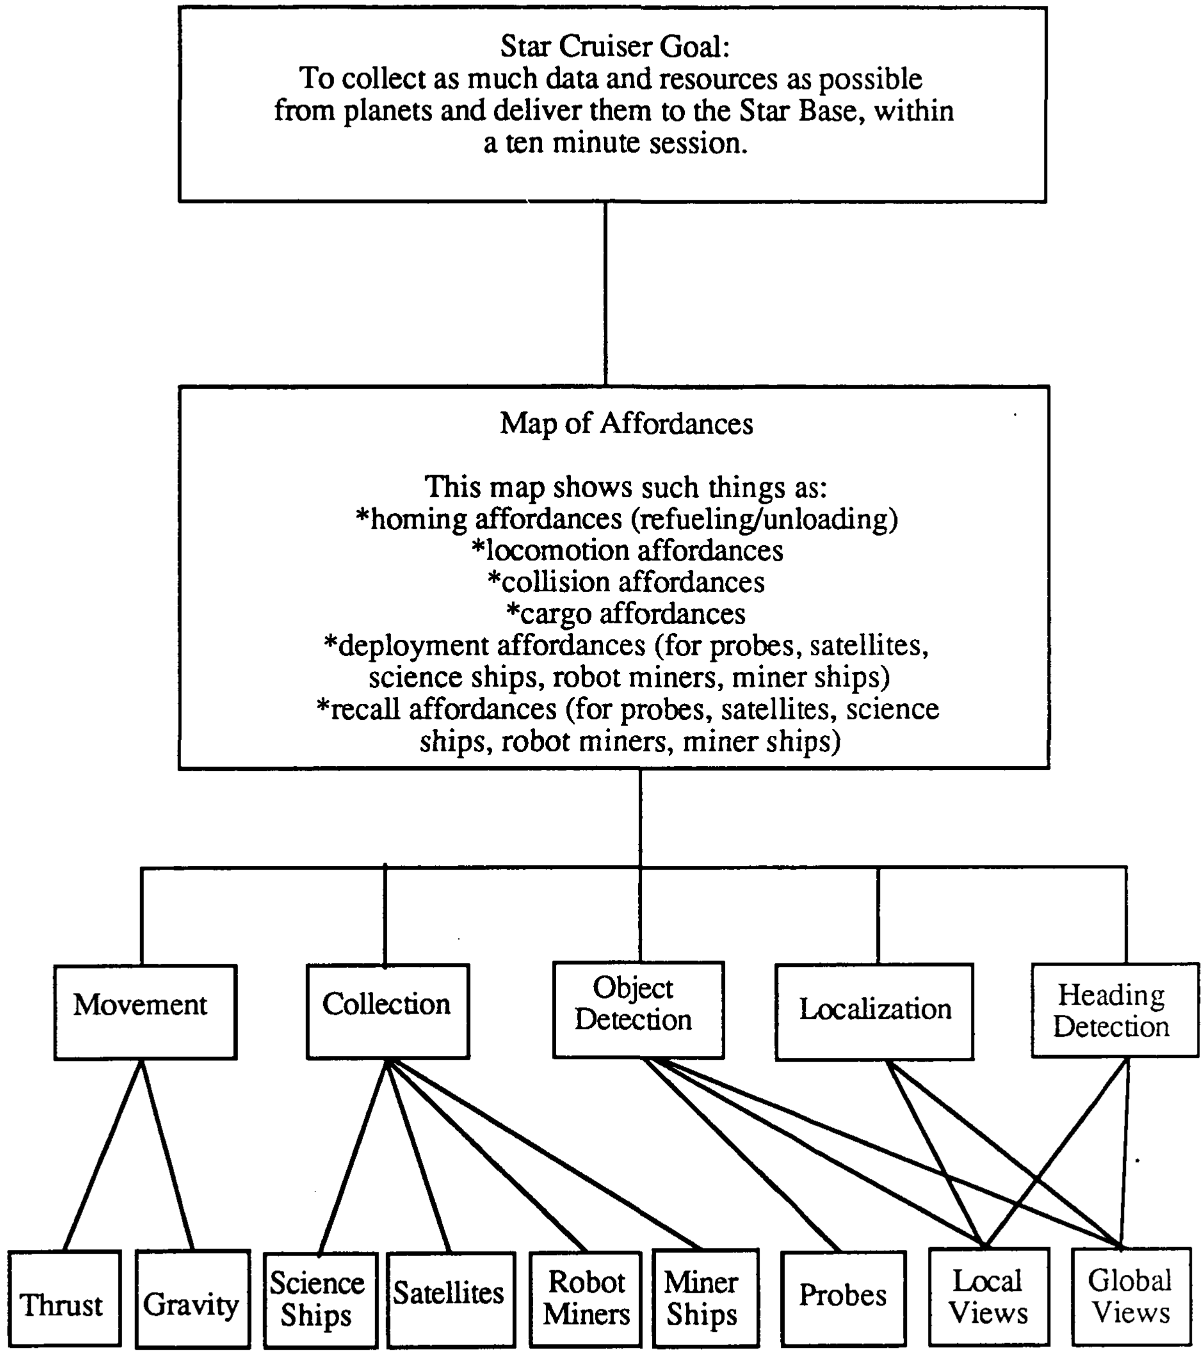

Figure 2 Preliminary Abstraction Hierarchy for Star Cruiser

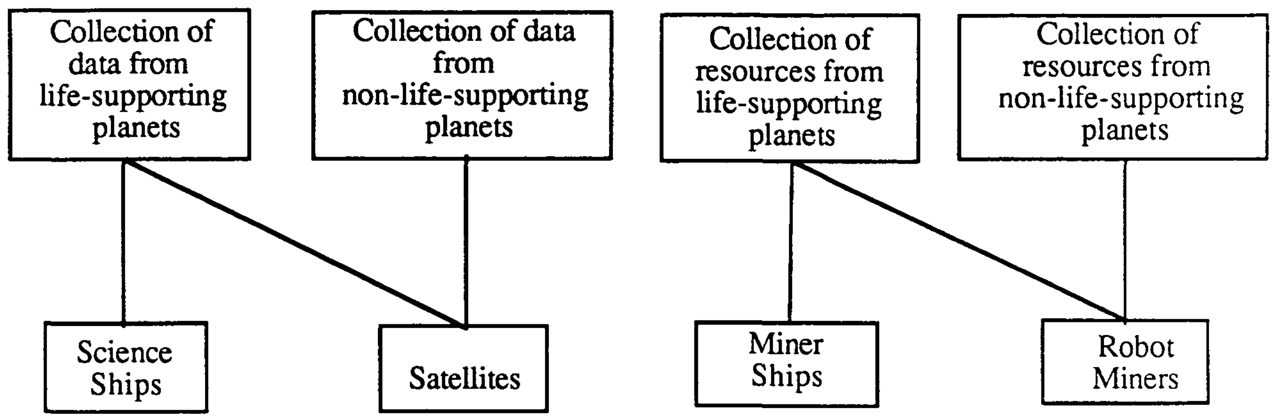

Figure 3 Decomposition of the Collection Function

Expected Contributions

The research proposed here hopes to integrate the independent research done on each of the two main decision-aiding technologies—intelligent decision-aiding and enhanced displays. It also hopes to determine which types of systems are better aided by which type of aid by carefully studying how domain characteristics interact with the characteristics and possible error modes associated with each of the two technologies. Accomplishing this task would greatly contribute to the areas of interface design and human-computer interaction, by providing guidelines for system design.

Additionally, by studying the interaction effects between operator skill level and type of decision aid, we hope to better understand the differences between expert and novice decision-making by discovering how each is better aided and under which conditions each is better aided. Finally, by studying the interaction effects between type of system trained on and type of decision aid, we hope to discover important guidelines and issues involved in the training of operators.

References

Allard, F. and N. Burnett. 1985. Skill in sport, in *Canadian Journal of Psychology, 39(2).*

Allard, F., S. Graham, and M.E. Paarsalu. 1980. Perception in sport: basketball, in *Journal of Sport Psychology, 2.*

Allard, F. and J.L. Starkes. 1980. Perception in sport: volleyball, in *Journal of Sport Psychology, 2.*

Chase, W.G. and H. A. Simon. 1973. Perception in chess, in *Cognitive Psychology, 4.*

Chi, M.T.H., P.J. Feltovich, and R. Glaser. 1981. Categorization and representation of physics problems by experts and novices, in *Cognitive Science, 5.*

Deakin, J.M. and F. Allard. 1991. Skilled memory in expert figure skaters, in *Memory and Cognition, 19.*

DeGroot, A.D. 1965. *Thought and Choice in Chess,* Mouton & Co.: The Hague.

Dreyfus, H.L., S.E. Dreyfus, and T. Athanasiou. 1986. Five steps from novice to expert, in *Mind over Machine.* The Free Press: New York.

Hammond, K.R., R.M. Hamm, J. Grassia, and T. Pearson. 1987. Direct comparison of the efficacy of intuitive and analytical cognition in expert judgement, in *IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-17, No. 5.*

Hetherington, E.M. and R.D. Parke. 1986. *Child Psychology,* McGraw-Hill: New York.

Hopkin, V.D. 1988. Air traffic control, in *Human Factors in Aviation.* Academic Press: New York.

Kirlik, A. 1990. *Star Cruiser,* a design specification document submitted to Sterling Software.

Klein, G.A. 1989. Recognition-primed decisions, in *Advances in Man-Machine Systems Research, Vol. 5,* W.B. Rouse (ed.). JAI Press: Greenwich, CT.

Larkin, J.H. and H.A. Simon. 1987. Why a diagram is (sometimes) worth ten thousand words, in *Cognitive Science, 11.*

Lesgold, A., H. Rubinson, P. Feltovich, R. Glaser, D. Klopfer, and Y. Wang. 1988. Expertise in a complex skill: diagnosing X-ray pictures, in *The Nature of Expertise,* Chi, Glaser, and Farr (eds.). Lawrence Erlbaum Associates: Hillsdale, NJ.

Norman, D.A. 1988. *The Psychology of Everyday Things.* Basic Books, Inc.: New York.

Rasmussen, J. 1986. *Information Processing and Human-Machine Interaction.* North-Holland: New York.

Soloway, E., B. Adelson, K. Ehrlich. 1988. Knowledge and processes in the comprehension of computer programs, in *The Nature of Expertise,* M.T.H. Chi, R. Glaser, and M.J. Farr (eds.). Lawrence Erlbaum Associates: Hillsdale, NJ.

Staszewski, J.J. 1988. Skilled memory and expert mental calculation, in *The Nature of Expertise,* M.T.H. Chi, R. Glaser, and M.J. Farr (eds.). Lawrence Erlbaum Associates: Hillsdale, NJ.

Suchman, L.A. 1987. *Plans and Situated Actions,* Cambridge University Press: Cambridge, MA.

Vicente, K.J. 1988. Adapting the memory recall paradigm to evaluate interfaces, in *Acta Psychologica, 69.*

Vicente, K.J. and J. Rasmussen. 1990a. Ecological interface design: theoretical foundations, unpublished manuscript.

Vicente, K.J. and J. Rasmussen. 1990b. The ecology of human-machine systems II: mediating "direct perception" in complex work domains, in *Ecological Psychology, 2(3).*

Woods, D.D. 1986. Paradigms for intelligent decision support, in *Intelligent Decision Support in Process Environments,* E. Hollnagel, et. al. (eds.). Springer-Verlag: Berlin.

Woods, D.D. and E.M. Roth. 1988. Cognitive systems engineering, in *Handbook of Human-Computer Interaction,* M. Helander (ed.). Elsevier Science Publishers

Appendix A

Hypotheses for the Star Cruiser Experiments

(Each letter here represents the performance of the corresponding subject group found in Figure 1. DE = display enhanced version. DA = intelligent decision-aiding version)

A<B: equal training (unaided), DE leads to better performance than original unaided

A<C: equal training (unaided), DA leads to better performance than original unaided

A<D: something learned during training (on DE) is good and transferred (both tested on unaided)

A<E: something about DE is better than unaided

A<F: either something about training w/DE or testing on DA leads to better performance than unaided

A<G: something learned during training (on DA) is good and transferred (both tested on unaided)

A<H: either something about training w/DA or testing on DE leads to better performance than unaided

A<I: something about DA is better than unaided

B>C: same training (unaided), DE leads to better performance than DA

B>D: DE leads to better performance (nothing learned or not enough to offset during training)

  B<D: more transferred during training w/DE

B<E: (want a little bit) but close to equal performance, diff. comes from training (equal testing—DE)

B F

B>G: DA leads to better performance (not enough transferred during training to offset)

  B<G: more transferred during training w/DA

B<H: something from DA transferred during training to improve performance on DE (better than unaided)

B>I: DE leads to better performance than DA w/o training on it

C<D: something during training w/DE transferred and led to better performance

C>E: DA leads to better performance than DE even w/o training

C<F: something during training w/DE leads to better performance on DA than training unaided

C>G: DA leads to better performance (not enough transferred during training)

  C<G: more transferred during training w/DA

C H

C<I: want a little better—shows better performance due to training w/DA (equal testing—DA)

D<E: otherwise DE during training transferred and accommodated (equal training—DE)

D<F: DA leads to better performance than unaided w/ equal training on DE

D>G: something during training w/DE transferred and improved performance on unaided more than training w/DA

D H

D<I: otherwise more transferred during training w/DE

E>F: otherwise equal training (on DE) but DA improves performance

E>G: otherwise more transferred during training w/DA than training and testing on DE

E>H: otherwise something during training w/DA transferred and led to improved performance on DE

E>I: something about DE better than DA

F G

F H

F<I: otherwise something about DE during training transferred and led to improved performance over DA

G<H: equal training (DA), DE better performance

G<I: equal training (DA), DA better performance than unaided

H<I: otherwise equal training (DA), DE better performance than DA

Appendix B

Listing of Star Cruiser Actions

The following is a listing of possible actions available on the Star Cruiser application. A total of five factors are discussed for each. The first paragraph under each action heading explains when the action **can** and **cannot** be performed or, in other words, when the program will allow the user to perform the action and when the user will not be allowed to do so. The second paragraph explains when the user **should** and **should not** perform the actions. This section is based on my own experience with the application and it details those times performing an action can be beneficial or when it can be detrimental to the user's performance. The third paragraph under each heading details what **perceptual support** exists for that action. It explains what support currently exists, whether it is satisfactory or not, and possible improvements. Once again, this is based on my experience with Star Cruiser and thus may differ from someone else's opinion.

There are several characteristics about Star Cruiser that one should keep in mind as s/he read through this listing. The first is that whenever one of Star Cruiser's tools needs to be deployed or recalled, the user must be viewing the map corresponding to Star Cruiser's location (i.e., in galaxy - global map; in solar system - local map). In addition, "movement" of Star Cruiser can only occur if it is not docked at Star Base or in an orbit. If it is, then Star Cruiser must first be taken out of orbit or pulled away from Star Base before it can travel freely.

Finally, it should be realized that many of Star Cruiser's movements are in preparation for the user to perform some other action. Most movements are the result of the user wanting either to deploy or recall one of Star Cruiser's tools or have Star Cruiser return to Star Base in order to refuel/unload cargo. Also, the choice (direction, speed) of movements may also depend on what information is obtained through viewing the global or local maps. It becomes apparent that, due to these factors, Star Cruiser's movement through the galaxy and solar systems is usually quite dependent on other actions that the user has just performed or wishes to do in the near future.

Deploy Probe

A user may deploy a probe anytime except when docked at the Star Base. In other words, the only time a probe can not be deployed is when the Star Cruiser is docked at the Star Base.

Deploying probes has no effect on points or fuel consumption and therefore can be done almost anytime the user wants to. It is advisable to deploy probes when the user has difficulty locating the 9th orbital in order for the Star Cruiser to orbit a sun. Deploying a probe is also useful when the user wishes to know the amount of data/resources in a particular solar system before visiting it. If the user has little difficulty in obtaining orbit with the Star Cruiser and/or doesn't need to know the amount of data/resources before visiting a particular solar system, then probe deployment has little value.

There is currently no perceptual support that informs the user when it is possible or best to deploy a probe. Since the deployment of a probe has no effect on fuel consumption or the user's score, there is no real need to inform the user when a probe should be or shouldn't be deployed. Blacking out the probes from the selection bar at the top of the screen can be of useful in preventing the user from trying to deploy a probe while docked at the Star Base.

Recall Probe

A probe in orbit around a sun may only be recalled by the user if the Star Cruiser is in orbit around that same sun and the user is viewing the system in local mode. If the user is viewing the galaxy, if the Star Cruiser is not in orbit in the same solar system as the probe, or if there are no probes in the solar system currently being visited by the Star Cruiser, then no probes can be recalled.

Probes only need to be recalled if they are to be used somewhere else in the galaxy. If that is the case, then it is suggested that the user recall a probe when all data/resources have been collected from the particular solar system and/or the user no longer needs assistance in identifying the 9th orbital. If the user still has trouble getting the Star Cruiser in orbit, then it is recommended that the user do not recall the probe.

If the program determines that the user would like to deploy another probe, but none are available, then it can highlight a probe that has already been deployed and is present in a system that contains no more data/resources. This would let the user know immediately the most ideal probes to recall. This, though, would be rather redundant in that the user can simply view the pie-chart present on the suns in global mode to determine which solar systems no longer contain any data/resources. Therefore, a significant amount of perceptual support already exists in helping the user to identify the probes which can or need to be recalled.

Deploy Satellite/Robot Miner

A Satellite/Robot Miner can be deployed only when the Star Cruiser is in orbit around a sun which contains planets and the user is currently viewing that particular solar system's local map. If the Star Cruiser is not in orbit or in a solar system or if that solar system doesn't contain any planets or if the user is not viewing the local map of the solar system containing the Star Cruiser, then deployment of a Satellite/Robot Miner is not possible.

A Satellite/Robot Miner should be deployed to any blue planets containing data/resources that need to be collected. They should also be deployed to any green planets that contain data/resources in order to prevent any Science Ships/Minerships from automatically collecting to much data/resources and thus preventing themselves from being recalled for fear of overloading the Star Cruiser with too much data/resources. The only time it isn't beneficial to deploy Satellites/Robot Miners is when there are no planets in the solar system which contain any data/resources.

No extensive perceptual support exists that helps the user decide whether or not to deploy a Satellite/Robot Miner and if so, which planet to deploy it to. The only clues that are present to the user are the Star Cruiser's gauges that relate, qualitatively, how much data/resources is currently on board the Star Cruiser. Because this action is one of the more important ones performed by the user, better support should be present. One possibility is to automatically highlight a Satellite/Robot Monitor (accompanied with a auditory signal) to signal to the user that one can be deployed. In addition, by highlighting a particular planet, the user would also know where best to deploy the Satellite/Robot Miner. It is questionable whether or not planets that contain too much data/resources for Star Cruiser to handle should be highlighted. The absence of any highlighted Satellites/Robot Miners would indicate that the user should not deploy any.

Recall Satellite/Robot Miner

When Star Cruiser is in orbit in a solar system where Satellites/Robot Miners are deployed to planets and the user is viewing the solar system in local mode, then those deployed Satellites/Robot Miners may be recalled. If the Star Cruiser is not in orbit in a solar system where Satellites/Robot

Miners have been deployed or if it is the galaxy or if the user is viewing the galaxy, then Satellites/Robot Miners may not be recalled.

Satellites/Robot Miners should be recalled when they are finished collecting the data/resources from their particular planets. Care must be taken not to recall them if they have collected so much data/resources that it would overload the Star Cruiser. Therefore, the user should recall the Satellites/Robot Miners before they complete their missions if, in not doing so, they run this risk.

As with deployment, the only support present to the user in making the decision when to recall a Satellite/Robot Miner are the Star Cruiser's gauges. Once again, the need is present for better perceptual support. Deployed Satellites/Robot Miners can be highlighted when the program feels it is best to recall them. This, of course, will depend on the amount of data/resources currently present aboard Star Cruiser and how much the Satellites/Robot Miners have and/or can collect at their planets. If they can collect all of the data/resources without resulting in an overload when recalled, then the program can highlight the Satellites/Robot Miners when they have completed their missions. If the planets contain too much data/resources, then the program can highlight them before they finish, thus informing the user that they need to be recalled as soon as possible. If the Satellites/Robot Miners have already collected too much for one reason or another, then the program will not highlight them until the Star Cruiser has unloaded it's current haul at Star Base and returned to the current solar system.

Deploy Science Ship/Minership

Science Ships/Minerships can be deployed under the same circumstances as Satellites/Robot Miners with the one exception that green planets must be present in the solar system since they may only be deployed to a planet which "supports life." If no green planets are present in the solar system, or if any of the other conditions similar to Satellites/Robot Miners are not met, then Science Ships/Minerships cannot be deployed.

Science Ships/Minerships should be deployed whenever green planets are present in the solar system and contain data/resources. They should, however, not be deployed if the total data/resources that will be collected by any one Science Ship/Minership will overload the Star Cruiser. Therefore, if this risk exists, then Science Ships/Minerships should not be deployed.

*** Refer to the discussion of perceptual support for Satellites/Robot Miners. The issues discussed there may also be applied to the Science Ships/Minerships. ***

Recall Science Ship/Minership

*** All issues discussed under <u>Recall Satellite/Robot Miner</u> may also be applied here. The one exception is in regard to highlighting the Science Ships/Minerships. Since these ships will move from green planet to green planet, collecting all available data/resources and because the possibility exists that these ships may collect so much data/resources that they would even overload an empty Star Cruiser, the program should inform the user when to deploy multiple Science Ships/ Minerships (by highlighting them) in order to prevent this. This prevention is accomplished by dividing the available amount of data/resources amongst various ships so that the smaller portions may still be loaded onto Star Cruiser. ***

Move Star Cruiser Into Orbit

Star Cruiser can be placed in orbit whenever it is present in a soar system. If it is moving about the galaxy, then Star Cruiser cannot be placed into orbit around any of the suns.

The user should place Star Cruiser into orbit around a sun whenever satellites, robot miners, science ships, and/or minerships are to be deployed in that particular soar system. Also, if any of them, along with probes, are to be recalled from that same system, then Star Cruiser must also be in orbit. It becomes unnecessary to place the Star Cruiser in orbit in a solar system if there are no planets present or if the need/desire to deploy or recall any ships or probes does not exist in that system.

The program already gives some hints to the user as to when it is necessary to obtain orbit with the Star Cruiser. The pie-charts (in local mode) representing the amount of data/resources available on a planet are an indication as to when ships should be deployed or recalled. These hints, however, do provide a direct mapping between the desired situation (deploy/recall ship) and the means with which to obtain the situation (put Star Cruiser in orbit). Therefore, more support is needed. One possibility is to highlight both Star Cruiser and where it should be (9th orbital). This, though, would have the drawback of showing the user exactly where the orbital is located thus making the action almost too simple to perform and also removing one of the probe's functions (identify 9th orbital). Displaying a message such as "Achieve Orbit" on the screen, which should be just as informative, would be a better option in that it would not simplify the task or remove any functions from the Star Cruiser's tools.

Move Star Cruiser Out Of Orbit

This action can only be performed if the Star Cruiser is already in orbit in some solar system and the user is currently viewing that same system.

The Star Cruiser should be moved out of orbit if the user has completed the task of either recalling as many deployments as desired or if the user wishes to exit the solar system for some reason such as moving to a new solar system or going to the Star Base. It is advisable. however, that Star Cruiser remain in orbit in order to recall as many deployments as possible as long as the risk of overloading on data/resources does not exist or there is no threat of running out of fuel.

No perceptual support exists that aids the user in determining when is the most opportune time to move out of orbit. None is really needed either. If enough support exists which informs the user of other actions to perform with Star Cruiser (i.e., dock at Star Base, recall satellite, etc.), then the user should know that in order to perform those tasks, Star Cruiser must or must not be in orbit. If the user does not know this though, a simple message can be used to provide instruction.

Move Star Cruiser Into Solar System

If Star Cruiser is moving about the galaxy, the user then has the option of moving it into a solar system. Star Cruiser cannot move directly from one solar system to another without first entering the galaxy. Nor can Star Cruiser enter a solar system if it is docked at Star Base. It must first pull away from Star Base, then it may enter a solar system.

It is beneficial to have Star Cruiser enter a solar system if that system contains any data/resources that the user wishes to collect. Thus, if the user wants to deploy any ships, then Star Cruiser must first enter the solar system before anything else can be done. This also holds true if the user wants to recall any ships. The user should try to prevent Star Cruiser from entering any solar systems if the current task is to get Star Cruiser to Star Base so it can dock. This is critical if Star Cruiser is low on fuel since it may not be able to reach Star Base if it keeps entering and exiting solar systems.

The only support in determining when to enter a solar system that the program provides the user with are the pie-charts that may be located on the suns in the global mode. Noticing whether or not a particular solar system contains any data/resources can help the user decide if it is worth entering. There may be many solar systems. however, that contain data/resources. Therefore, the program should also suggest to the user which particular solar system Star Cruiser should enter. This can be done by printing the message "Enter Highlighted Solar System" along with highlighting a particular sun.

Move Star Cruiser Into Galaxy

If Star Cruiser is in some solar system, but is not in orbit, then the user may move Star Cruiser directly into the galaxy without having to perform any other intermediate tasks such as moving Star Cruiser out of orbit first.

When the user has completely loaded up Star Cruiser with data/resources and has already taken it out of orbit, then Star Cruiser should be moved into the galaxy so that it can make its way to Star Base. In addition, whenever Star Cruiser no longer needs to remain in a solar system, it should be moved into the galaxy so that it may travel to another system or to Star Base. Star Cruiser should more than likely not move into the galaxy if there still remains more data/resources that can be collected without causing an overload of Star Cruiser and if Star Cruiser is not at risk of running out of fuel.

There is no direct assistance provided to the user that says when Star Cruiser should exit the solar system and enter the galaxy. However, the pie-charts depicting the available data/resources shown on the planets, or their absence, should help the user determine whether or not it is worth staying in the solar system. In addition, the fuel gauge and Star Cruiser's gauges showing its remaining capacity for data/resources also help the user decide if Star Cruiser need to move into the galaxy so it can go and dock at Star Base. This, though, is generally enough support. Other assistance such as informing the user to dock at Star Base should provide further help in determining when to enter the galaxy.

Dock Star Cruiser At Star Base

The user may only dock Star Cruiser at Star Base if Star Cruiser is present in the galaxy and the user is viewing the global map. If Star Cruiser is in any solar system, then it cannot dock at Star Base.

Star Cruiser should dock at Star Base whenever it cannot carry any more data/resources or whenever it is about to run out of fuel. If Star Base is nearby, though, and Star Cruiser still can carry more data/resources without becoming overloaded and still has plenty of fuel, it is sometimes good strategy to dock at Star Base to unload the cargo and refuel. This proves beneficial when it comes time to have Star Cruiser journey to those solar systems which are far from Star Base. Star Cruiser should not be forced to dock at Star Base if it is not necessary if the base is far away. Since these missions have a time limit, actions of this nature will only waste that time.

The fuel gauge and Star Cruiser's data/resource gauges help the user determine when it is necessary to have Star Cruiser dock at Star Base. This should generally be enough support. During high workload situations, however, the user may forget to check these gauges. Therefore, as a safety precaution, it is probably wise to display some message informing the user that Star Cruiser better dock at Star Base. This would appear only under "must"-situations. The user should be allowed to determine whether or not Star Cruiser should dock without the use of any other additional information besides the gauges.

Have Star Cruiser Leave Star Base

The user can have Star Cruiser leave Star Base right after it has docked there.

Star Cruiser should be made to leave Star Base right after docking since no other actions can be performed until it has done so. The only time it would not be necessary to leave Star Base is when all data/resources have been collected, thus ending the scenario.

No perceptual support informing the user to pull Star Cruiser away from Star Base is needed. The fact that nothing else can be accomplished until Star Cruiser's departure provides enough of a forcing function to remind the user to do so.

View Galaxy (Star Cruiser In Solar System)

The galaxy may only be viewed if Star Cruiser is in orbit in some solar system. The only other time that the global map is viewed is when Star Cruiser is traveling through the galaxy itself.

This action's purpose is merely to gather information about various states of the system. Such items that may be checked by the user include the collection status of the total amount of data/resources in a different solar system; the distance from the current solar system to Star Base; or the proximity/location of other solar systems to the current system. This action is useless if the user does not desire any such information.

Because this action merely provides information to the user (it does not alter the system states in anyway), no perceptual support is required. If the user desires some piece of information that can only be gathered through viewing the global map, then the user will select that option. If it cannot be selected, then the user will realize that Star Cruiser is not in orbit and that it may be simpler to just move Star Cruiser into the galaxy. Since there is no way of determining which information the user would like to have access to, it is difficult to have the program support the decision to view the global map.

View Solar System (Star Cruiser In Galaxy)

If a probe has been deployed to a particular solar system or if Star Cruiser has previously visited it, then that system may be viewed while Star Cruiser is traveling around the galaxy. The only other method for viewing a solar system is to have Star Cruiser enter the system.

This action merely provides information to the user. Such information may include the number of planets in a particular solar system or the collection status of deployments in that system. If no information is desired about a particular system, then the user need not perform this task.

No perceptual support for this action is needed. Since no changes are being made to the system states and it is difficult to know exactly when the user desires information, let alone what kind, it would be almost pointless to try to provide any support for deciding when to perform this action.

REQUIREMENTS FOR PSYCHOLOGICAL MODELS TO SUPPORT DESIGN: TOWARDS ECOLOGICAL TASK ANALYSIS

Alex Kirlik
Georgia Institute of Technology

Abstract

Cognitive engineering is largely concerned with creating environmental designs to support skillful and effective human activity. The goal of this chapter is to propose a set of necessary conditions for psychological models capable of supporting this enterprise. An analysis of the psychological nature of the design product is used to identify a set of constraints that models must meet if they can usefully guide design. It is concluded that cognitive engineering requires models with resources for describing the integrated human-environment system, and that these models must be capable of describing the activities underlying fluent and effective interaction. These features are required in order to be able to predict the cognitive activity that will be required given various design concepts, and to design systems that promote the acquistion of fluent, skilled behavior. These necessary conditions suggest that an ecological approach can provide valuable resources for psychological modeling to support design. Relying heavily on concepts from Brunswik's and Gibson's ecological theories, *ecological task analysis* is proposed as a framework in which to predict the types of cognitive activity required to achieve productive behavior, and to suggest how interfaces can be manipulated to alleviate certain types of cognitive demands. The framework is described in general terms, and illustrated with an example from our previous research on modeling skilled human-environment interaction.

Note: This paper is a draft of a chapter for a forthcoming book, The Ecology of Human-Machine Systems, J. M. Flach, P. Hancock, J. Caird, and K. J. Vicente (Eds.), Hillsdale NJ: Erlbaum. Please do not cite or distribute. Comments are welcome to: kirlik@chmsr.gatech.edu

This research was supported by NASA Ames grant NAG2-656 to the Georgia Institute of Technology, Robert J. Shively, technical monitor.

INTRODUCTION

Modern psychology judges its progress and products by a variety of criteria. Reviewing a number of paradigms in current cognitive psychology, Claxton (1988) suggests that the research community gives no less that thirteen answers to the question: "How do you tell a good cognitive theory when you see one?" Each of the thirteen criteria he mentions (e.g., experimental, computational, evolutionary) has enough adherents so that research programs are judged successful even if their products meet perhaps only one of these standards of merit. Research activity in current cognitive science thus resembles a massively parallel search, where most of Claxton's thirteen criteria for scientific success are suspended on any one search path so that individual research efforts can proceed unencumbered by a diverse set of otherwise paralyzing constraints. For example, in certain paradigms computational realization is the primary concern, mathematical formalization the major constraint in others, and in still others a necessary condition for a theoretical model may be a demonstration that the proposed cognitive mechanisms and processes could have emerged through human development or evolution. The eventual success of this divide and conquer venture, of course, hinges not so much on whether each of the many paradigms meets its own goals, but rather, on whether we are somehow able to integrate the resulting array of research products into useful and coherent theory.

It is natural to wonder whether we have decomposed our research efforts in a way that will allow for eventual theoretical unification. One primary concern is whether the many research paradigms that comprise cognitive science are moving along diverging or converging paths. Perhaps this is a question best left for time to decide. I am concerned, however, that although strict and dogmatic adherence to a single scientific criterion may lead to individually successful hillclimbing, when considered overall we may find we have all climbed different hills, and if anything, actually increased the difficulty of the journeys between us. A coherent, useful cognitive theory will have to meet a large number of constraints. Rarely, however, do good solutions to problems which involve meeting multiple constraints emerge by decomposing the problem via the constraints themselves. Knowing the least expensive restaurant in town, the one with the best food, and the one with the healthiest menu is not particularly helpful in allowing one to find a good square meal at a fair price.

The purpose of this chapter is to identify a set of necessary conditions for psychological models capable of supporting the design of environments to promote skillful and effective human activity, i.e., *cognitive engineering* (Fischoff, Slovic, and Lichtenstein, 1978; Norman, 1986; Rasmussen, 1986; Woods and Roth, 1988). This effort is motivated by my own limited success in attempting to apply the products of cognitive science to cognitive engineering. My experiences have led be to believe that the central problem that needs to be overcome to make the products of cognitive science more relevant to design is identifying a more productive set of dimensions along which modeling efforts can be decomposed. We simply must decompose the complex problem of cognitive modeling in order to make any headway. However, for cognitive engineering at least, the decomposition must be derived from an overall framework capable of ensuring that the resulting research products can be reassembled into a coherent theory useful for design.

A description of a solution to any problem, even if expressed only as a set of necessary conditions, plays a crucial role in formulating a problem decomposition strategy capable of ensuring that the subproblem solutions can be effectively integrated. We have to know where we are going if we want to get there. In terms of the previous analogy, we have to know that our goal is a good square meal at a fair price in order to determine how to decompose the problem of finding an appropriate restaurant. A necessary step toward a more applicable cognitive science, therefore, is a statement of the the set of constraints that must be met if a psychological model is to support design. There may be no good reason to expect that the set of constraints that must be satisfied to support design are identical to the set of constraints cognitive scientists normally use to guide their scientific explorations. In fact, I will suggest below that the necessary conditions for an acceptable psychological model in cognitive science are quite different than the necessary conditions for a psychological model capable of guiding design. Many of the difficulties involved with trying to apply cognitive science modeling arise out of this mismatch. Cognitive science has simply decomposed its central problem in a manner that is very unfortunate for the cognitive engineer.

As Carroll (1991) has noted in regard to the failure of psychology to meaningfully contribute to understanding the problem of human-computer interaction (HCI), the realization that the products of a basic science do not provide effective resources for application can provide important lessons for the basic science itself (also see Flach, 1990b; Gibson, 1967/1982; Neisser, 1976). The solution to the problem of creating a scientific basis for cognitive engineering is not merely one of improving the designer's access to research findings (e.g., Meister, 1989), moving research into naturalistic or operational contexts (e.g., Klein, 1989), or improving generalizablity from experimental results (e.g., Hammond, Hamm, and Grassia, 1986), although each of these goals is surely important. Rather, I am convinced that the solution must lie in a reformulation of the questions posed by the basic psychological research itself: a reformulation driven by an understanding of the psychological nature of the design product and the knowledge that is required in order to create it.

THE PSYCHOLOGICAL NATURE OF THE DESIGN PRODUCT

A standard modeling approach in cognitive psychology is to hold a task environment relatively fixed and to create a description of the cognitive activities underlying a person's behavior in that environment. The designer, on the other hand, is faced with the opposite challenge of creating an environment to elicit a desired behavior, with the ultimate goal being the creation of a design that is maximally consistent with the principles underlying how people skillfully and effectively interact with the world. In problem solving terms, the solution space for the scientist is the set of plausible cognitive theories while the solution space for the designer is the set of technologically feasible environments. We can thus characterize the scientist's problem as a search among possible cognitive "solutions" to a given task, and the designer's problem as a search among possible environments to obtain a given cognitive solution. These are symmetrical psychological problems of comparable subtlety and difficulty, requiring equally sophisticated empirical and theoretical methods. In this sense a theoretical/applied dichotomy does not appear to be a faithful way of portraying the difference between the practices of cognitive psychology and environmental design.

One reason, though, for the perpetuation of the theoretical/applied distinction is a lack of appreciation for the psychological nature of the design product. While the scientist creates theories of cognitive function, it is assumed that the designer creates not theories but merely environments: a mix of hardware and software that is best conceived in technological rather than in psychological terms. But this perspective is based on an overly restrictive view of what the environment is, from the standpoint of understanding human behavior. Although a design product may be implemented in hardware and software, this is the wrong level at which to view the relevant features of that product, just as it would be wrong to look for the relevant features of a psychological theory in the software or ink in which it is realized.

As Carroll and Campbell (1989) have noted, each design product is actually an instantiation of the designer's theories of how the environment influences how people behave, think, and skillfully perform, however rudimentary and fragmentary these theories may be. Although the (good) psychological theory is only implicit in the design of a (good) VCR interface, for example, it is nonetheless real in exactly the same sense that the physical theories and electrical engineering principles implicit in the design of the inner workings of the VCR are real. If you want to understand or predict the functionality of the VCR, you had better know the operative physical theories underlying its design. Similarly, if you want to understand or predict human interaction with the VCR, you had better know the psychological theory underlying its design as well. While the VCR can of course be looked at as an assemblage of physical matter, this is the wrong level at which to look for the relevant structure of the machine, either for understanding electromechanical function or for understanding user interaction. Both the electrical engineer and cognitive engineer structure the physical matter using organizing principles derived from theories within their own disciplines. It is the adequacy of these theories, rather than any facts solely about the physical form of the machine, that determine whether the VCR will play, and whether the user can play it.

Each instance of human interaction with any artifact is thus a psychological experiment testing the assumptions embodied in the environmental design (cf. Wise's (1985) construal of an architectural design in terms of scientific hypotheses). Although it may be fashionable within the cognitive engineering community to bemoan how little guidance modern psychology provides the designer, the psychological nature of the design product is inescapable. The correct response to the current and unfortunate lack of applicable psychological research is not to attempt to do psychology-free design (since this is impossible -- the design will not be apsychological but instead reflect the designer's "folk" psychological theory), but rather to ask what kinds of psychological models are needed to support cognitive engineering, and to begin the long range empirical and theoretical work necessary to realize them.

This is not to imply, of course, that cognitive engineers should wait for a more applicable psychology to emerge before making design commitments (this would be hopelessly naive), or should turn all attention to modeling and away from the design of prototypes and expanding technological opportunities. It may be, as Braitenberg (1984) has suggested, that human capacities for synthesis far exceed capacities for analysis, in the sense that our creative products reflect a degree of implicit or tacit knowledge that is far more elaborate and rich than the knowledge we can explicitly state and formalize. The direct manipulation interface was not deduced in any interesting sense from psychological theory -- in fact, a case could be made that we still have no unified psychological theory that would predict the profound superiority of direct manipulation over command interfaces for various tasks. It just so happened that in this case the designer had an implicit understanding of how people naturally interact with the world that was closer to the truth than any explicit and formal psychological theory available at that time. While cognitive psychologists may be able to identify why command-line interfaces are inefficient in various ways, it was the designers and not psychological research that pointed toward environments that support more efficient interaction, a finding that should probably inform psychological research itself.

There is a catch, however, to this design-as-research strategy. Assuming a particular prototype of a design concept is successful, any useful generalizations that emerge from creating the prototype will be at the level of the psychological assumptions underlying the design, rather than at the level of the particular technologies used to implement the design. To return to a point made earlier, the hardware and software implementation is the wrong level to look for the relevant features of the design product. Especially in HCI, a vast amount of research effort has been expended trying to answer questions comparing various interface technologies, for example, design options such as scrolling windows, hypermedia, etc. This research is of dubious value (see also Vicente, this volume), since the "it depends" answers produced by such efforts will only lead to a neverending series of technology-specific design principles, rather than a stable and generative theoretical account of human-environment interaction that can guide design in novel situations.

To profit from the design-as-research strategy, then, it is incumbent upon the researcher to make explicit the psychological assumptions that contributed most to the success of the prototype system. A successful system demonstrates nothing other than its own success, unless the possibly implicit psychological theory underlying the design is articulated. Although forcing the researcher to articulate the theoretical assumptions prior to environmental design (as would be demanded by traditional experimental psychology methodology) may actually impede progress -- synthesis may be more efficient than analysis -- the hope for generalizable conclusions from such demonstrations surely rides on whether the researcher can subsequently identify the psychological hypotheses that were validated by the success of the prototype. There is probably no alternative to traditional experimental methodology for this purpose. A research program using the design-as-research strategy must include both an initial synthesis phase followed by an analysis phase where the implicit psychological theory guiding synthesis is made explicit, tested, and communicated.

MODELING TO SUPPORT DESIGN

I have argued that a good psychological theory is a necessary aid to design by discussing the essential psychological nature of the design product, and also by showing that the problem faced by the designer is not one of mere application but is instead itself a theoretical problem comparable to that faced by the scientist. A search for environments to promote a particular mode of cognitive activity and behavior (the designer's task) is no more an applied endeavor than is a search for accounts of cognitive activity and behavior that are promoted by particular environments (the scientist's task). There is, however, an important difference between these two problems. The designer and scientist search in opposite directions; one reasons over environmental models while the other reasons over cognitive models. It should be expected, therefore, that different types of heuristic guidance will be necessary to direct search in the two cases. As a result, the theories that best provide heuristic guidance to the scientist will have different properties than the theories that best provide heuristic guidance to the cognitive engineer. Let us examine how the theories that would best support reasoning over environments might differ from the theories that would best support reasoning across cognitive activities.

Much of our current understanding of cognitive-level human-environment interaction consists of a set of somewhat independent <environment-process-behavior> triples, each of which provides a psychological model of how a person might achieve a particular behavior in a specified environment. When the difficult but important job of integrating this knowledge into coherent theory is attempted, these efforts typically focus on integrating across the process and behavior dimensions rather than across the environment dimension. The result is that understanding the environmental contribution to behavior is a largely ignored component in the theoretical unification. One approach, for example, to achieving theoretical unification of this set of triples is to integrate across the behavior dimension. The results here are powerful, typically hybrid "cognitive architectures" (Card and Newell, 1989). These general purpose cognitive frameworks have the processing resources to produce a wide variety of behaviors, from simple motor responses to complex problem solving and planning. Yet another approach is to integrate primarily across the process dimension in an attempt to show that the functionality of a wide variety of existing models can be subsumed under a single process modeling formalism. Cognitive models demonstrating how symbolic processing techniques can be implemented using neural network or connectionist formalisms are good examples of partial theoretical unification along the process dimension.

Theoretical integrations along the environment dimension, however, are hardly ever attempted but are critically needed to support the cross-environmental reasoning inherent in design. It should not come as a surprise that most cognitive psychologists are not overly concerned with this type of theoretical unification, since an acceptable scientific product is a model of behavior in a specified environment, and rarely is reasoning backwards from cognitive theory to environment required. Except perhaps in experimental design itself, rarely is the cognitive psychologist forced to reason across environments in order to activate specified cognitive modes. Significant exceptions (i.e., attempts at theoretical integration across environmental influences on cognition and behavior) are Rasmussen's (1986) theory of multi-level environmental representation as reflected in the "abstraction hierarchy," and Hammond, Hamm, Grassia, and Pearson's (1987) efforts to obtain a rich enough set of environmental and task descriptors so that the cognitive mode underlying judgment behavior (e.g., analytical, intuitive) can be predicted and promoted through environmental manipulation. Only a unified theory of the environmental influences on cognition can guide the designer's search for environments to activate specified cognitive processes.

But what would a unified theory of the environment look like, and what types of guidance would it provide? What would be integrated in such a theory would be the diverse set of knowledge of what the psychologically relevant aspects of the environment might be, for the purpose of trying to understand or predict human behavior and performance. One must, for example, determine when it will be appropriate to understand the environment in terms of stimuli and reinforcements as in behaviorism; cues, criterion and feedback as in models of judgment; options, chance nodes, choice nodes and probability distributions as in decision theory; initial states, goal states and operators as in Newell and Simon's (1972) problem solving theory; affordances or constraints on action as in Gibson's (1979) ecological theory, system state variables and differential equations as in manual control theory, and so on.

Each of the above forms of environmental description has its place. No single representation of environmental structure will do justice to understanding the many different forms of cognition and behavior observed in complex human-environment interaction (Rasmussen, 1986). The reflection of environmental structure in behavior is manifest in various ways, and each way is suggestive of a different model that best describes the structure of the environment to which productive behavior must be sensitive. A large part of design activity, in fact, can be viewed as the selection of appropriate environmental descriptions. In some cases, the cognitive engineer faces the problem of selecting an environmental description for an existing candidate design that will assist in predicting the cognitive activity and behavior the environmental design will promote. In other cases in which the cognitive engineer can operate earlier in the design cycle, the central problem will be to create a design concept, expressed as an environmental model, that promotes a specified mode of cognitive activity maximally consistent with the demands of a task. Guidance for both of these cognitive engineering activities can only come from theoretical frameworks that support the designer's reasoning over alternative environmental models.

Psychology may already have the rudiments of such a theory, but perhaps oddly, this knowledge is expressed not so much in existing models of cognitive activity, but rather in the process of designing experiments capable of successfully activating those cognitive activities for scientific study. That is, it is the often tacit and unformalized knowledge guiding experimental design in studies of cognition that approaches the type of understanding needed to reason effectively over environmental models. When successful, the knowledge underlying the experimenter's ability to promote a particular mode of cognitive activity is quite similar to the type of knowledge necessary to guide system design. Much has been made of the inability of basic experimental psychology research to guide design (e.g., Rouse, 1987, Meister, 1984), but perhaps the fundamental difficulty is that the knowledge the designer needs goes beyond the experimental findings, it may approach the knowledge needed to have actually *designed* these experiments.

ISSUES IN ENVIRONMENTAL MODELING

Cognitive engineering thus demands techniques for environmental modeling with a strong theoretical basis, and the resulting environmental models must be as explicit, formal, and precise as the models used to describe internal cognitive activity. A cognitive psychology capable of predicting environmentally situated behavior and of supporting design will therefore have to be concerned as much with the environment as with internal cognitive activity (e.g., Brunswik, 1952; Gibson, 1979; Anderson, 1991). When one looks at the types of models produced by current cognitive psychology, however, rarely does the environmental model receive close to the amount of attention as does the internal cognitive model. There are at least three reasons for this state of affairs.

First, experimental psychologists often feel the need to simplify their environments for the purposes of control, and are thus able to get by with highly simplistic and impoverished environmental models (compare the length of the stimulus description -- the environmental model -- with the length of the description of the internal psychological model in most papers in the cognitive experimental literature). Second, and especially in research within the cognitive science orientation, often no distinction is even made between the description of the external environment and the subject's internal representation of the environment. While it may indeed be the case that interesting questions can be answered using such an approach (e.g., differences in the types of internal representations used by expert and novice problem solvers), these accounts start so far downstream that they fail to capture any influences of the external problem representation upon the efficiency of problem solving activity (but see Larkin and Simon, 1987). Finally, rarely is it the case that researchers working within an established paradigm are forced to reason across widely varying environmental conditions, with the result that assumptions about environmental descriptions can remain implicit within a given research program. There may be no pressure to unconfound the environmental from the internal constraints on cognition and behavior when environmental manipulations are made over a very narrow range.

Thus, the open problem for cognitive engineering is to determine under what environmental conditions various cognitive activities will be activated, and required, for effective task performance. To evaluate a candidate design, the issue is not only to understand cognitive processes such as problem solving, decision making, and working memory, but also to determine *what* problems will have to be solved, *what* decisions will have to be made, and *what* working memory demands will be, given various design concepts for a particular task.

In the following I will discuss two types of constraints on acceptable psychological models that arise due to the need to represent both internal and external influences on cognition and behavior.

The first set of constraints are structural. I will argue that the a model's structure must be capable of representing both cognitive and environmental organization in a single, unified format; i.e., that the appropriate unit of analysis and modeling must be the human-environment system, rather than the human alone. The second set of constraints concern the content of acceptable models. I will suggest that cognitive engineering is most in need of environmental models that assist in understanding fluent, skilled human interaction with the world, rather than environmental models that rationalize detached intellectual activity. In most cases the design goal is (or should be) to create a design which promotes fluent and skilled activity, rather than a design which promotes cognitively-intensive control of behavior . We require environmental models that capture the features of environments that promote effective, skilled performance in order to define a design target and also to identify the causes of error-prone cognitive activity in current systems.

Modeling the Integrated Human-Environment System

One of the earliest attempts to model human-machine interaction concerned manual control behavior, such as steering a car or flying an aircraft. Engineers familiar with the design of electromechanical feedback control systems turned their attention to modeling the human as a feedback control system in order to assess human capabilities and limits so that vehicles could be designed so that control demands were within these limits. Control theory has a well specified language for environmental modeling. The controlled "plant" (airplane, automobile) can be described in terms of a *transfer function* that relates system inputs (steering adjustments) to system outputs (heading). The human as a feedback controller can be described in similar terms. In this case the input might be the heading of the automobile and the output would be a steering command. As Flach (1990a) notes, the goal in this endeavor was to discover the human transfer function; i.e., a description of the function relating stimuli to response during manual control behavior. At this schematic level of description, much current psychological modeling shares this goal of finding invariance at the organismic level, rather than at the level of the organism-environment system.

These engineers were in for a rude awakening however, as empirical results indicated that the there was no single human transfer function. Rather, the human transfer function appeared to adjust to changes in the dynamics of the controlled system. As Birmingham and Taylor (1954) noted, the ability of the human to adjust to the environmental transfer function was so great as to suggest "that '<u>the</u> human transfer function' is a scientific will-o'-the-wisp which can lure the control system designer into a fruitless and interminable quest." (p. 1752) Subsequent modeling attempts (McRuer and Jex, 1967) were only successful once the search for invariance in behavior shifted to the level of the human-machine system, rather than in human behavior alone. The crossover model of human manual control behavior developed by McRuer and his colleagues is a statement of behavioral invariance at the level of the human-environment system.

Why should this finding concerning human perceptual-motor behavior inform our discussion of cognitive-level human-environment interaction? The answer is that there appears to be little reason to expect that cognitive-level behavior will be any *less* adaptive to environmental structure than is perceptual-motor activity. In fact, there are a variety of reasons to believe that just the opposite is that case; i.e., that human cognitive interaction with the world is even less constrained, and thus more flexible, than is perceptual-motor interaction. Note also that the correct response to this situation, and the one pursued by these manual control researchers, is to describe both human and environment as an integrated unit, and to use this unified human-environment model as a tool in the search for behavioral invariance. Pursuing such a strategy requires formalisms capable of expressing both internal and external constraints on behavior in the *same* language, such as the transfer function representations used to model both the manual controller and the controlled system.

Although rarely used, this approach has been successfully applied to understanding cognitive-level behavior. The Lens Model framework for the description and analysis of human judgement (Brunswik, 1952; 1956; Hammond, 1955) is a unified description of both the human judge and the environment. As such, it has been a fruitful tool in understanding both the environmental and cognitive constraints on judgment abilities, and has been enlightening as to a number consistencies in judgment performance that would likely not have surfaced without some mechanism for partialing out the environmental contribution to behavior (e.g., see Brehmer and Joyce, 1988). The recent book by Anderson (1990) describing the Rational Analysis framework also represents a step in the direction toward integrated human-environment system modeling. This framework provides resources to address the question of how both internal "computational" constraints and external environmental structure combine to determine the processes that will be engaged to perform a particular task.

When we turn to the problem of understanding the kinds of environmentally "situated" (Suchman, 1987; Whiteside, Bennett, and Holtzblatt, 1988) activity typical of behavior in modern human-machine systems, it is clear that much work remains to be done before an integrated human-environment modeling approach will be possible. However, the kinds of psychological descriptions already being proposed for describing dynamic human interaction with technological systems indicate a clear shift toward understanding how both environmental and cognitive structure contribute mutually to the production of skilled behavior.

In a pair of penetrating analyses of the cognitive-level ecology of human-machine systems, Hutchins (1987; 1991) has suggested that human cognition and behavior cannot be understood apart from the external devices in the environment that have been designed to perform cognitive functions. In ship navigation (Hutchins, 1987), for example, human interaction with notepads, checklists, and calculators can sometimes be used in lieu of memorial, procedural, and computational operations; and in modern aircraft (Hutchins, 1991), much of the cognitive burden for memory of intended and current speeds has been allocated to external memory structures within the cockpit. In such environments, the entire cognitive function is distributed across both person and environment. It is not surprising, then, that understanding these integrated systems requires describing both internal and external cognitive functions in mutually compatible terms. We have come full circle: the computer metaphor that gave rise to a description of human cognition in terms of information processing has been turned back upon the world, as seen in the description of the environment as external memories, external problem representations, and the like.

The importance of these environmental aids to thought and behavior cannot be underestimated. Much of modern psychological research paints a rather dismal picture of human cognitive abilities and limitations, leaving some of us in a state of wonder over how it can even be possible for human cognition to have resulted in its modern achievements. But rarely does even the scientist work in isolation from external cognitive tools, as Donald (1991) has noted.

> For example, there is no internal wiring schema to support the kind of synthesis made possible by a scientific diagram; the synthesis is *out there*, in the diagram itself. The theoretician depends heavily upon a huge variety of external cognitive props -- mathematical notations, curves, plots, histograms, analog measurements, and technical jargon -- to arrive at a theory. Without these things, thoughts of this kind would simply not be possible, because the end-state or "conclusion" reached by the mind is driven directly by the external representation itself. The locus of a process like theoretical synthesis would thus be difficult to attribute to any single part of the internal-external network that makes up such a system. (pp. 378-379).

The same comments would also apply, and perhaps in even greater force, to understanding the mechanisms underlying skilled activity, such as flying an airplane, driving a car, or performing the many routine tasks we find in daily life. Skilled activity is often accompanied by a heightened level of intimacy with the world rather than by increased detachment, an observation that leads to the hypothesis that intensive exploitation of environmental structure plays a key role in productive behavior. As Norman (in preparation) has noted:

> With a disembodied intellect, isolated from the world, intelligent behavior requires a considerable amount of knowledge, lots of deep planning and decision making,

> and efficient memory storage and retrieval. When the intellect is tightly coupled to the world, decision making and action can take place within the context established by the physical environment, where the structures can often aid as a distributed intelligence, taking some of the memory and computational burden off the human. (Chapter 10, p. 6)

The human-environment system must serve as the unit of analysis and modeling to allow the internal cognitive activity necessary for productive behavior to be predicted as a function of environmental design, and also to identify how necessary cognitive activity can be engineered through environmental manipulation.

The Need for Models of Fluent Interaction with the World

Since the goal of the cognitive engineer is often (but not always) to create an environmental design that promotes fluent and effective skilled behavior, the features of environments that support fluent as opposed to cognitively intensive behavior need to be identified and described. In many existing human-machine systems, the reason that complex cognitive processing is necessary for effective performance is that the environments in which the operators work are quite unlike those environments in which human psychological abilities evolved. As a result, the acquisition of fluent modes of behavior is impeded and the end state is one of only partially effective adaptation. The problem of "situation awareness" (e.g., Sarter and Woods, 1991) in the modern commercial aircraft cockpit is a prime example. Edwards (1988) has gone so far as to describe the cockpit as an "opaque veil," and Bohlman (1979) suggests that the difficulty of maintaining an active understanding of the aircraft and airspace from cockpit displays is so great that it is appropriate to speak of crews as constructing "theories" of their situations.

As one who tries to make a living constructing theories, I find it most unsettling to think that theoretical abilities are sometimes necessary to ensure safe flight. What kind of psychological theory would provide the most leverage for remedying the situation awareness problem? Because cognitively intensive activities such as inductive inference, hypothesis generation, and mental modeling are observed in current systems, it seems only natural that better accounts of activities such as these are the key to enhancing interaction. Such accounts could presumably guide the design of aids to assist flight crews in their theoretical tasks, or the design of training methods to make crews better theoretical thinkers. It is natural to view such attempts with suspicion, however, since problem solving aids have the potential to create their own set of human-machine interaction problems (Woods and Roth, 1988), and training to make people more "rational" problem solvers or decision makers has yet to be proven effective.

The alternative solution, of course, is to design environments more consistently with the principles underlying skilled, dynamic human interaction with the world. Pursuing this strategy, however, requires techniques for environmental modeling capable of representing the features of task environments that both promote and inhibit the acquisition of fluent modes of behavior. Only models of productive, skilled behavior can provide the resources for a task-analytic approach capable of identifying features of an environmental design that are inconsistent with the principles underlying skilled activity. The problem of identifying demands for complex cognitive activities such as problem solving, planning, decision making posed by a given environmental design most requires models of skilled, fluent behavior, *not* models of problem solving, planning, or decision making activities.

Why is this the case? Due to their roots in either economic theory or artificial intelligence, rational action models such as those mentioned above are more concerned with sufficiency considerations than they are with necessity considerations. The great appeal of such models is their ability to describe and often prescribe behavior in a huge variety of situations. Nearly any, and perhaps all, behavior can be rationalized as being the result of some cognitively-intensive process such as search through a problem space, hypothetico-deductive inference, or the comparative evaluation of options with respect to a goal structure or utility function. No empirical evidence could ever be brought to bear on limiting the *sufficiency* of these rational methods for action selection. However, identifying when these sorts of complex cognitive activities will actually be *necessary* for successful performance requires models capable of indicating when such activities are *not* necessary.

My observations of skilled human behavior in complex systems have led me to the working hypothesis that cognitively-intensive methods for action selection are used only as a last resort; i.e., when effective perception-action solutions are not readily available. Predicting cognitive demands thus requires modeling approaches capable of defining when effective perception-action solutions will not be available, and this knowledge can only be provided by a theory of perception-action skill. I realize that the claim that skilled performers will typically opt for perception-action solutions to cognitive tasks may strike the cognitive psychologist as being counter-intuitive. However, intuitions based mainly on laboratory findings may be skewed by the fact that experiments on cognition are typically carefully designed to preclude the availability of perception-action shortcuts for meeting task demands. Although my own intuitions are largely based on observations of behavior in operational settings, even in the laboratory I am continually amazed at the cleverness of subjects who are able to short-circuit demands for complex cognitive activity by cuing off the whir of a disk drive or an aberration in the graphics software. I have ceased to be suprised and frustrated by such cleverness, and have begun to view the tendency toward the perceptual selection of action as a fundamental aspect of skilled behavior. There is no doubt, however, that more empirical research is needed to clarify this issue. However, the necessary experiments must provide rich enough environmental conditions and enough practice time so that both cognitively-intensive and perception-action task solutions are made available. Such laboratory experiments are rarely conducted.

These "experiments," however, are performed every day in both complex human-machine systems and in more everyday work settings. As an example, over the past two years I have made fairly extensive observations of the behavior of short-order cooks working busy rush periods at a local area grill. My interest in this behavioral situation arose because skilled performance in this setting appeared to possess many of the same properties I have observed in my more limited observational studies in complex operational settings, and also because 24 hour access to this environment can be readily secured for the price of a cup of coffee. And by making a well timed food order or by initiating conversation with the cook, one even has an (albeit limited) capacity for intervention and control over task demands.

In the environment I have studied, the cook uses an assortment of automated devices such as fryers and ovens, combined with substantial manual activity at the grill, to coordinate the preparation of the many items within each order, while preparing multiple orders simultaneously. Describe this task in any formalism for rational action and the task demands appear overwhelming. Observe this type of skilled human-environment interaction, though, and I believe the following will be apparent. First, there is an intensive degree of intimacy in the cook's perception-action interaction with the environment. The cook maintains tight perceptual contact with the world and always seems to be taking some sort of action. Rarely if ever does the cook appear to engage in detached, contemplative cognitive activity. Task demands are uncertain and arrive dynamically, and ongoing behavior must be sensitive to a number of unpredictable events.

What allows this perceptually intensive mode of interaction to be productive? Note that the cook's environment is highly structured, but nearly all of this structure is *visible*. The most efficient "problem representation" for the cook to use is an external one: the grill area itself. Action selection based on the external environment has considerable economies as compared to action selection based upon internal representations of the environment. The environment considered as a problem representation serves as an external memory capable of being perceptually accessed, updates itself automatically and in parallel, serves as an external memory store, is internally consistent, and is always veridical (also see Reitman, Nado, and Wilcox, 1978). The world takes care of its own "truth maintenance."

When uncertainties do occur using such an external representation (i.e., perceptually available information underspecifies constraints on activity), these uncertainties can often be resolved through perception-action rather than accessing stored knowledge. How well cooked is the underside of a steak? Flip it and see. And the cook not only uses the structure already present in the environment, he or she can dynamically *create* structure in order to make perception-action solutions available and thereby reduce cognitive burdens. For example, the cook may organize the placement of meats in order of doneness, may lay out dishes or plates to serve as a temporary external memory of orders to be prepared (also see Beach, 1988), and may even generate new information "displays" by introducing constraint in the controlled environment causing a hidden variable to covary with a visible one. For example, the cook may adopt the strategy of continually flipping meats so that the doneness of the top side can always be used as a reliable indicator of the doneness of the underside. In a very real sense, the cook is both performer and on-line interface designer.

Skilled human-environment interaction of this type is thus both a response to environmental structure as well as a source of environmental structure to be subsequently exploited. The environmental structure created by the cook's own "tricks" and routinized strategies plays a role similar to the structure created through the environmental design process itself in promoting cognitive efficiencies. The former structuring merely happens "on-line" and is thus short-lived, while the latter happened during the design of the grill area and is thus reflected in the static and permanent organization of the design. But both forms of structure, whether contributed at one point in time by the designer or continually by the cook, result in cognitive economies through the enablement of perception-action solutions to the task. For example, the external memories and displays dynamically created by the cook play a similar cognitive role to the timing mechanisms used in toasters and ovens to offload memory demands to the world. Because of the possibility of self-produced environmental structure, the acquisition of such situated skills will always resist faithful description solely in terms of the development of more efficient internal mechanisms for processing a fixed set of environmental information. A model of skill acquisition in dynamic human-environment interaction would also have to describe how the actor's external environment becomes increasingly structured by activity itself, and thus increasingly informative to the actor, over the course of skill development.

Much of the responsibility for dynamic human-environment interaction lies in the perception-action mechanisms at the interface between the performer and the world. The development of skilled, dynamic interaction relies upon abilities to exploit environmental structure to obtain perception-action solutions to tasks, and where none naturally exist, to create additional environmental structure in such a way as to enable perception-action solutions. If such structure is not provided by the designer, the performer will seek to create it through activity that introduces new forms of structure. The productivity of this mode of behavior requires the availability of sources of information to specify the environmental constraints to which behavior must be sensitive in order to be effective, and the availability of actions capable of both changing the environment and of creating additional sources of information to further enable the perceptual guidance of activity. These are, I believe, features common to nearly all environments in which the acquisition of fluent, dynamic interaction is observed. They are also features lacking in the many technological environments of interest to cognitive engineering, due largely to interfaces that highly restrict perception-action access to the controlled system. The absence of such features is one major cause of the difficulty of acquiring skills in such systems, and the reason that the end state of learning is often one of only partially effective adaptation.

TOWARD AN ECOLOGICAL PERSPECTIVE

The previous discussion has centered on identifying a number of necessary conditions for psychological models to support environmental design. It is time now to turn toward outlining a methodological strategy with the potential to address the some of the gaps in our knowledge discussed above. Many of the necessary features for psychological models that have been identified are suggestive of the possibility that an ecological approach to human-environment interaction may yield fruitful tools for cognitive engineering. The ecological approach was pioneered by Brunswik's (1952) and Gibson's (1966, 1979) theories of how knowledge of environmental structure can provide important constraints on psychological explanations. In particular, Brunswik's emphasis on taking the human-environment system as the unit of analysis and modeling, perhaps best represented in the Lens model framework (Brunswik, 1952; 1956; Hammond, 1955), and Gibson's focus on how fluent interaction can be described as perceptual specification of environmental constraints on activity (1979), blend nicely with the claims that cognitive engineering is most in need of models of skilled interaction with the world, and models which take the human-environment system as the unit of analysis.

In the following I will take some initial steps toward identifying opportunities the ecological approach might offer for cognitive modeling to support design. However, and for readers already familiar with Gibson's views especially, it is important to first discuss what an ecological approach to cognitive engineering does *not* require. First, it does not require that we conceive of all human-environment interaction as purely perceptually guided activity. Direct perceptual guidance of action, as discussed by Gibson, might surely be possible although it is likely that it is specific to those information-rich environments in which perception evolved or to artifactual environments designed to mimic such environments. There is no reason to expect that evolution anticipated the modern aircraft cockpit or the word processor. In such environments the need for post-perceptual processes such as problem solving and decision making is quite likely. The ecological and information processing approaches need not always be considered to be at odds, but may instead both contribute to a more complete understanding of human-environment interaction.

Second, the adoption of an ecological approach does not necessarily imply a commitment to studying fluent behavior in the natural environment. Gibson rallied against the use of abstract information displays for the study of visual perception; the types of displays often found in existing human-machine systems. But in a larger sense, Gibson, like Brunswik before him, was arguing for using environmental conditions as the basis of scientific study that are representative of the conditions in which a target behavior of interest occurs. And for better or worse, a cockpit or control room looks much more like a laboratory than it does the natural terrestrial environment. These are the target environments of interest to the cognitive engineer. For this reason, these environments, or carefully made abstractions of them, are the places where the ecological approach to cognitive engineering should be carried out.

Resources for Cognitive Modeling

Brunswik (1952) offered the Lens model as a description of how the human and environment could be described in an integrated fashion, using the principle of *parallel concepts* (e.g., see Hammond, Steward, Brehmer, and Steinmann, 1975). As shown in Figure 1, the Lens model is a symmetrical framework which represents how both environmental and cognitive structure mutually contribute to judgment performance. The organism has available a set of cues (x_i's) which bear specified relations ($r_{e,i}$'s) to an environmental criterion to be judged (e.g., a medical diagnosis). The relations between the the cues and the criterion may take various forms and vary in *ecological validity*. Similarly, the ways in which the organism's makes use of the cues ($r_{s,i}$'s) to arrive at a judgment may take various forms and vary in *cue utilization*. The framework is an expression of the principle of parallel concepts in that each concept on one side of the model has a counterpart on the other side. This framework has a number of attractive properties that result from representing the organism and environment in compatible terms.

Insert Figure 1 about here

Perhaps most importantly, the Lens model framework allows the modeler to measure the degree to which the environmental structure which relates the cues to the criterion is reflected in the manner in which the cues are cognitively structured to produce a judgment. High levels of achievement are an indication of a highly *adaptive* cognitive organization; i.e., a cognitive strategy that mirrors the environmental structure to which behavior must be sensitive. As many of the previous comments in this chapter suggest, some sort of adaptivity-oriented view of cognitive activity is likely to be required in order to understand skilled human-environment interaction as well. In addition, the Lens model allows one to localize the causes of less than fully productive behavior to either the environmental structure, the cognitive structure, or both. Weather forcasters, for example, frequently err in their predictions, but only an analysis of both the ecological validity of the cue structure and their policies for cue utilization can yield an understanding of the reasons for these errors. We often have similar interests in the design and analysis of human-machine systems. Did a particular error result from an operator making incorrect usage of displayed information, or was the error the result of a potentially perfectly adapted operator confronted with not fully diagnostic information? Quite different types of remedial action can (and should) result depending upon the answers to questions such as these.

Thus, the Lens model framework offers a good starting point for developing an approach for representing skilled human-environment interaction. However, my previous comments suggest that we not only require models that take the human-environment system as the unit of analysis, we also need models capable of representing fluent, skilled behavior in order to identify demands for more complex cognitive activity. From this perspective, the Lens model has two important deficiencies. First, action itself is not explicitly represented. The Lens model is a epistemological framework for the purpose of modeling judgments about the state of the world, not to represent how actions are selected. A first step toward applying the Lens model to action selection would be to allow an action opportunity itself to serve as the criterion to be judged. However, this interpretation can give rise to a number of conceptual difficulties, and we must be careful in how we go about formulating this interpretation in order to keep distinct the environmental model (facts about the world), and the cognitive model (facts about the performer). Second, as suggested by Hammond et al. (1975), judgment is a "cognitive activity of last resort" (p. 272). Judgmental abilities will only be called upon when the available information only probabilistically specifies the criterion, and actions capable of manipulating environmental variables to gain more diagnostic information are not available. Note that these environmental properties are exactly those features of task environments I have previously described as being the major impediments to the development of fluent, perceptually guided interaction. Like formalisms for rational action, then, it may be quite possible to interpret the skilled selection of action within the Lens model framework, however, such a model is not likely to capture those special features of task environments that allow for the acquistion of fluent human-environment interaction.

The need to explicitly represent fluent interaction, as well as those features of the world that promote it, suggest that we must consider the problem of how the environment of the skilled performer should be described. Gibson, with his ecological physics and theory of affordances (1979), proposed an action-oriented environmental description in order to understand how perception may orient behavior to environmental opportunities for action. An action-oriented approach results in a description of the environment in terms of the opportunities for action it presents the performer. The resulting description can be called an *affordance space*, akin to the decision space descriptions resulting from decision theory, the problem space descriptions resulting from problem solving theory, or the cue space descriptions resulting from theories of probabilistic judgment. In most cases, an affordance space will be a dynamic description of the environment, as both the environmental structure and the performer's resources for action will change over time and a dynamic affordance structure will result.

Note that creating an affordance space environmental description does not commit one to any particular position concerning how affordances may be detected to guide activity. Gibson was most concerned with those situations in which perceptual information is available to specify affordances, and in such cases interaction can be described as the perceptual detection of information capable of orienting behavior to action opportunities. However, in other situations perception-action access to the environment may be restricted or impoverished, or information other than that specifying the immediately present affordance structure must be taken into account for behavior to be productive. In such cases, performers may have to engage in more elaborate cognitive activity in order to detect the affordance structure, or to combine information specifying affordances with other information in order to select actions. Regardless of the type of either perceptual or post-perceptual activity required to orient behavior to an affordance structure, a description of the world in terms of affordances is still a valuable tool in understanding how environmental structure is reflected in cognition and behavior.

For our purposes, the concept of an affordance space is especially important because it provides resources to compensate for the two deficiencies of the Lens model identified above. First, an affordance space can in some cases play the role of the criterion in the Lens model, thus shifting the emphasis from passive judgment to the identification and selection of opportunities for action. Second, the possibility that perceptual information is capable of fully specifying the affordance space (i.e., direct perception) suggests that we must relax the *a priori* assumption underlying the

Lens model that the available information is only probabilistically related to the criterion. The question of determining the relationship between the available information and the environmental affordance structure is an empirical one. The ecological task analysis framework presented below is an attempt to integrate Brunswik's Lens model and Gibson's affordance theory into a unified framework for modeling skilled human-environment interaction.

The Framework for Ecological Task Analysis

Integrating concepts from Brunswik's Lens model and Gibson's affordance theory results in the *ecological task analysis* framework depicted in Figure 2. Like the Lens model, the proposed framework is a symmetrical arrangement which represents the integrated human-environment system. Not only does ecological task analysis exploit the principle of parallel concepts to capture certain symmetries between cognitive and environmental structure, the proposed framework also uses a principle of parallel concepts to suggest certain symmetries between perception and action. This latter symmetry is evident in the relationship between the upper model of environmental perceptual structure, and the lower model of environmental action structure.

Insert Figure 2 about here

Ecological task analysis begins with the creation of two complementary descriptions of the *surface* structure of the environment. As in the Lens model framework, we rely upon a distinction between environmental surface structure which exists at the interface between the performer and the world, and the environmental *depth* structure which exists remotely, behind the surface structure, so to speak. The two descriptions of surface structure, shown in the middle of the diagram, are models of environmental perceptual structure and environmental action structure. The description of surface perceptual structure in ecological task analysis plays a similar role to the cue description in the Lens model. The model of surface perceptual structure describes the environmentally available information. In human-machine systems, one can think of surface perceptual structure as the information available from interface displays about the state of the controlled system. Unlike the Lens model framework which captures only the surface perceptual structure of the environment, however, ecological task analysis also requires a description of the surface *action structure* of the environment. In a human-machine system, one can think of surface action structure as the actions made available by interface controls.

In the Lens model framework, the relation between surface perceptual structure and depth structure is the relation between the readily available information and the environmental state a person is attempting to judge. Similarly, in the proposed framework the relation between surface action structure and depth structure is the relation between the readily available actions and the environmental change the person is attempting to effect. On the perceptual side, the surface/depth distinction reflects the difference between given and inferred. On the action side, the surface/depth distinction reflects the difference between readily available actions and intended actions. The model of environmental depth structure on the left side of Figure 2 is a description of these potentially covert relationships, and the model of cognitive structure on the right side of the figure is a description of how environmental depth structure is reflected in cognition, and ultimately, behavior.

The Principle of Parallel Concepts: Perceptual and Action Structure

The description of environmental surface structure in terms of both perception and action is necessary to capture how the skilled perceptual guidance of activity is possible. The possibility for perceptual guidance of activity relies upon the availability of perceptual information capable of fully specifying environmental affordances. In terms of the proposed framework, perceptual guidance of action is not to be described as the perceptual detection of surface information to infer some covert or remote depth property of the environment. Rather, the perceptual guidance of action is to be described as the use of one form of surface structure, namely perceptual, to specify another form, namely, the surface action structure. As will be described in detail below, the ecological task analysis process proceeds by examining the models of perceptual and action surface structure to identify the congruence between these two forms of environmental description. Matches between these two models indicate opporunities for fluent, perceptually guided activity. Mismatches between these two models are indicative of demands for cognitive activity to overcome the perceptual non-specification of action. Various forms of mismatch are possible: each form of mismatch is suggestive of a different type of necessary cognitive activity, and a different type of remedial interface design solution.

Before describing the analysis process, the symmetrical nature of the models of perceptual and action structure must be discussed. The model of surface perceptual structure is an environmental description using a performer s perceptual capacities as a frame of reference in which the envirommental structure is described. This description is relational in the sense that the resulting environmental model reflects both the perceptual capacities of a performer and the environmental structure. It is easy to overlook that such descriptions are actually relational in nature, as we often speak as if a perceptually-generated differentiation of the environment is purely a function of the environmental structure and not a function of perceptual capacities. However, the fact that some forms of environmental structure are seen as objects while others are not, or the fact that some aspects of the environment are seen as being blue while others are red, are as much facts about the perceptual system as they are of the environment. A perceptually-oriented environmental model is a relational construct created by using perceptual capacities as a frame of reference in which environmental structure is measured and described.

The second model of environmental surface structure required for ecological task analysis is an action-oriented description of the environment. The model of surface action structure is an environmental description using a performer s action capacities as a frame of reference in which the environmental structure is described. This description is relational in the sense that the resulting environmental model reflects both the action capacities of a performer and the environmental structure. This action-oriented environmental model represents the world in terms of its opportunities for action. This environmental model thus generates a differentiation of the world in terms of the degree to which various spatiotemporal environmental regions are consistent in various degrees with the taking of various actions. An action-oriented environmental model is a relational construct created by using action capacities as a frame of reference in which environmental structure is measured and described.

The symmetrical nature of the perceptually-oriented and action-oriented environmental descriptions required for ecological task analysis should be apparent. These two environmental models differ only in that one is relativized to the functionality of the performer s input mechanisms while the other is relativized to the functionality of the performer s output mechanisms. The two models thus reflect two different primitive differentiations of the environment, one generated by using perceptual capacities to understand how the world is carved up with respect to perception, and the other generated by using action capacities to understand how the world is carved up with respect to action. As Barwise and Perry (1983, p. 11) have suggested,

> The emphasis is on how the organism differentiates its environment, on the sorts of uniformities it recognizes across situations. Different organisms can rip the same reality apart in different ways, ways that are appropriate to their own needs, *their own perceptual abilities and their own capacities for action.* This interdependence between the structure the environment displays to the organism and the structure of the organism with respect to the environment is extremely important. For while reality is there, independent of the organism s individuative activity, the structure it displays to an organism reflects properties of the organism itself. (My emphasis)

Neither the perceptually-oriented environmental description nor the action-oriented environmental description results in a more primitive, privileged, or objective ontological picture of the world. However, the claim is sometimes made that an action-oriented description of the world in terms of affordances is scientifically illegitimate because it is relativized to the actor, and is thus in some sense subjective. Note that the correct response to this possible criticism is not to adopt the heroic position that affordances are in some sense independent of the performer: they are not, they arise from using the performer s action capacities as a frame of reference for environmental description. Rather, to counter this argument one must merely emphasize that the supposedly scientifically legitimate perceptually-oriented environmental descriptions which portray a world of objects and properties are just as relativized to the capacities of the performer as are action-oriented environmental descriptions. One merely takes perceptual functionality as the frame of reference for environmental description, while the other takes action functionality as the frame of reference.

The Process of Ecological Task Analysis

I shall call a process whereby environmental models of surface perceptual and action structure are created and mismatches between the environmental differentiations represented in the two models are identified and described, an *ecological task analysis* of a human-environment system. What I believe to be the central contribution of an ecological approach to cognitive modeling can now be stated quite simply. A preliminary ecological task analysis of a human-environment system is required to identify the degree to which an interface (natural or artificial) between the human and environment is consistent with the principles underlying fluent interaction, and by doing so such an analysis helps specify *what* cognitive processes will be necessary for effective behavior. An ecological task analysis is thus similar in spirit to Marr s (1982) computational-level theory which attempts to define the necessary functionality of vision models, and Anderson s (1990) rational analysis which attempts to define the necessary functionality of a variety of cognitive models. Before outlining the analysis process, a few comments comparing the goals of ecological task analysis and Anderson s approach in particular may be valuable.

The goal of ecological task analysis is to define the necessary functionality of any cognitive processes that may be required to support effective human-environment interaction. Like Anderson s rational analysis, then, the present approach leaves aside the question of *how* any necessary cognitive activities might be carried out. In this much the two approaches share a common perspective. However, ecological task analysis takes human-environment interaction to be its primary concern, whereas the current formulation of rational analysis is concerned with cognitive activities such as memory, categorization, and problem solving. But as Anderson himself has noted, all these cognitive abilities are useless if they do not in some way serve the goal of action selection, and for Anderson, action selection is to always to be understood as the result of a problem solving or decision making exercise (p. 192). Perhaps it doesn t much matter what words Anderson uses to describe the processes underlying action selection, but what does matter is the nature of the environmental models that result from such a choice (decision trees and problem state spaces).

But as our previous comments suggest, all action selection, whether resulting from cognitively intensive processes or from more efficient perception and action, can be rationalized into one of these frameworks for explaining productive behavior. Modeling to support design, however, must focus not on the sufficiency of these frameworks but rather on limiting the conditions of their necessity. Ecological task analysis, therefore, starts not by assuming action selection is governed in any particular manner, but rather has as one if its goals to define what sorts of governing mechanisms will be necessary for modeling any instance of human-environment interaction. Ecological task analysis attempts to meet this goal through the use of environmental models that allow for a description of action selection in terms of the perceptual guidance of activity. Only when an examination of these environmental models indicates mismatches in perception-action environmental structure will ecological task analysis result in a construal of action selection in terms of decision making, problem solving, or any other cognitively intensive activity.

The Principle of Parallel Concepts: Environmental Depth and Cognitive Structure

As mentioned above, ecological task analysis begins with the creation of the models of perceptual surface structure and action surface structure which provide the bridge between the human and the environment. The analysis process proceeds by then examining the congruence between these two models in terms of the manner in which they differentiate the environment with respect to perception and action. The process of examining congruence will be discussed in detail below. The result of this exercise, however, is the specification of a model of the environmental depth structure, and a complementary model of internal cognitive structure and process. It is very important to note that ecological task analysis requires that modeling the environmental depth structure comes *after* creating the models of perception and action structure. That is, we do not begin by assuming that the task environment possesses a certain intrinsic depth structure (e.g., problem state space, decision tree, linear cue-criterion function), and thereby, a corresponding structure to cognitive processes (e.g., heuristic search, comparative evaluation of alternatives, linear cue combination rules). Rather, we let the examination of the surface perceptual and action structure of the environment guide the selection of the model of depth structure, and thereby the corresponding structure for cognitive processes. Herein lies the major difference between the proposed approach and many current modeling approaches in cognitive science: considerations of perception-action functionality define the necessary functionality for cognitive processes, rather than defining perception-action functionality by an *a priori* cognitive model of action selection.

I have tried to make the case for sequencing the analysis in this fashion at a variety of previous points in this chapter. We are obviously looking for a model of environmental depth structure that corresponds to, and helps make sense of, the cognitive activities that will actually be engaged to serve action selection. Which cognitive activities will be necessary for productive behavior, however, can only be determined by a detailed analysis of the degree to which perceptually information is available to specify productive action. Different interface (control-display) designs for the same task environment can differ radically in terms of the cognitive demands they make upon the performer, with the result that the cognitive model that provides the best description of the processes underlying interaction with one interface may be quite unlike the cognitive model that best describes the processes underlying interaction with another interface. As a result, different environmental models will be needed in the two cases to describe the different ways in which environmental structure is reflected in cognition and behavior in the two cases. Note, however, that like the Lens model framework, once a model of environmental depth structure and a corresponding model of internal cognitive processes are selected, the congruence between the environmental and cognitive models can be examined in order to identify any ways in which cognitive limitations or biases place constraints upon productive behavior.

This is obviously a highly schematic description of the models necessary for ecological task analysis. Yet I wonder if it is possible to get more precise about the content of these descriptions without doing a potential injustice to the richness of the perceptual and action structure of any realistically complex behavioral situation. One great allure of environmental models that rationalize action selection as decision making or problem solving is that relatively low-dimensional environmental descriptions can be used. Such low dimensional representations are advantageous in that they can be easily applied across a variety of contexts and are thus suggestive of how the psychological processes underlying action selection might be organized in a context-free, general purpose format. Describing the world in terms of the interaction of raw environmental structure with perceptual and action capacities, on the other hand, has the potential to create environmental models of almost unlimited dimension. Nevertheless, my own view is that any method capable of identifying the mechanisms underlying skilled human interaction in a setting of any reasonable complexity requires models capable of preserving many of the fine details of the environmental structure. The richness of the world s perceptual and action structure, seemingly necessary for the fluent operation of the perception-action system, severely overtaxes our highly limited cognitive-linguistic resources for environmental description. I doubt that at the current time a task analysis technique capable of guiding human-environment interaction modeling can possibly be much more than a charge to the modeler to undertake the long and arduous process of identifying and describing the potentially overwhelmingly rich interface between the performer and the world.

For this reason, the following description of the analysis process will consist of an abstract discussion of the possible results of an ecological task analysis, along with a concrete example of how such an analysis can be performed for a particular behavioral situation. Figure 3 depicts four possible results of an ecological task analysis in terms of the congruence of the resulting perceptually-oriented and action-oriented environmental models. The grid lines in each of the schematic environmental models indicate the manner in which the environment is spatiotemporally differentiated with respect to either perceptual capacities or action capacities. The four cases will be described separately.

Insert Figure 3 about here

Case I: Perceptual overspecification of action

In the first case shown in Figure 3, the perceived environment is over-differentiated with respect to the environmental differentiation in terms of opportunities for productive action. Many different perceptually distinct situations all point to a single opportunity for productive action. Object or configural displays are one type of design solution available for coping with perceptual overspecification of action. Object displays are an attempt to reduce the dimensionality of the perceptual space so it becomes aligned with the lower dimensional action space. These displays perform this function by organizing the originally over-differentiated perceptual information in such a way that perceptually salient relational features emerge that are differentiated in a manner identical to the differentiation reflected in the action space. When a display based solution is not used, however, the performer will have to develop some ability to overcome perceptual overspecification. Perceptual pattern recognition is one process that could potentially result in an alignment of the perception and action spaces, although some naturally occurring relational properties must be perceptually available to enable this solution. When the possibility for pattern recognition is neither naturally supported nor supported through configural display design, it is likely that a significant amount of categorical or instance-based learning may be required in order to identify consistencies in the mapping from perception to action. Note, however, that there is still the possibility of fully productive performance in all these cases since perception merely over-specifies action, it does not misspecify action, as in Case IV below.

Case II: Perceptual underspecification of action

Here the perceived environment is underdifferentiated with respect to the environmental differentiation in terms of productive action. There is simply not enough perceptually available information in order to uniquely specify the appropriate action alternative. One likely cause of perceptual underspecification of action is that there is hidden-state information in the environment; i.e., the performer must know something about the previous history of the environment, keep this information in memory, and then integrate this memorial information with the perceptual information in order to uniquely specify the appropriate action alternative. Building memory into the displayed interface using trend or historical displays is one strategy for aiding performance in such cases. Another cause of perceptual underspecification is that information about future state, rather than past state, is not perceptually available. Perception identifies a number of action candidates, but the selection of the appropriate action requires knowledge of the downstream effects of an action, and these effects are not perceptually apparent. Predictor displays or fast-time simulations are two approaches for making this information available to the performer. Unaided performance, however, will require considerable learning before skill can be acquired. Internalization of environmental dynamics in the form of an internal model may be necessary in order to gain access to past or future state information. The problem solving models of Newell and Simon can be viewed as descriptions of the cognitive activities that may be necessary when the downstream effects of an action must be taken into consideration. Exploratory behavior, or physically trying out solutions is a method available for taking into account information about future state without the use of an internal model, although this form of activity may not always be possible. A form of perceptual learning that is available to overcome initial perceptual underspecification is perceptual differentiation (e.g., E.J. Gibson, 1969). Here, the perceptual capacities of the performer change in order to increase the dimensionality of the perceptual space so as to bring the perceptual differentiation of the environment into alignment with the action-oriented environmental differentiation. Perceptual learning of this type, however, requires (perhaps initially subtle) dimensions of stimulation to which perception can eventually become sensitive.

Case III: Perceptual specification of action

In this situation fluent performance can be expected to develop without significant cognitive demands or conceptual learning. The performer s pre-established perceptual competencies provide the ability for unique specification of productive action. No rule-based information integration is required. Neither is information about the history or future of the environment necessary. Therefore, no internalization of environmental dynamics is required to provide these forms of information. If it were the case that the designer could always be certain that the perceptually-oriented and action-oriented environmental models of the human-environment system were correct, interfaces that support the perceptual specification of action would be our undeniable design target. However, if unanticipated changes occur in either the diagnosticity of the displayed information, the functionality of the interface controls, or the environmental dynamics, a fluent, informationally encapsulated perception-action mode of control may carry on without the performer paying heed to these environmental disturbances.

Case IV: Perceptual misspecification of action

In the final case the mapping between the perceptually available information and productive action is unruly. Behavior might well be productive in this situation, but not because the currently perceived situation is particularly informative. Models of behavior in such situations typically endow the performer with a considerable amount of knowledge to overcome perceptual misspecification, or else give up hope for a deterministic account of action selection and instead opt for finding invariance in aggregate performance through the construction of probabilistic cognitive and environmental models. In fact, the Lens model of human judgement can be considered to be a special case of the ecological task analysis framework under the assumptions that perception missspecifies action (judgment), and that the environmental depth structure can be described with cue-cue and cue-criterion correlations indicating the covert relationships among these variables. I will have little to say about this case because, frankly, it is something of a catch-all. However, I do think it is important to emphasize that when the modeler finds what appears to be a case of perceptual misspecification, but yet observes productive action selection based upon perceptual information, it is likely that the performer knows some things that the modeler does not. I suspect that such cases are more frequent than we may care to admit. The long history of findings on the context-sensitivity of reasoning (e.g., Wason and Johnson-Laird, 1972; Johnson-Laird, 1975), decision making (Kahneman and Tversky, 1979; Tversky and Kahneman, 1981), and problem solving (Kotovsky, Hayes and Simon, 1985) all demonstrate that people pay considerably more attention to the concrete presentation of a problem situation than do many abstract cognitive models. There is, of course, really no environmental stuff that is context as opposed to relevant structure. Context is always defined with respect to a model; it is simply those aspects of the environment that a given model fails to represent, and as a result, those aspects that are rendered incapable of producing behavioral variance. Findings which demonstrate the intensive context-sensitivity of cognition and behavior can be seen to be, in part, a reflection of the fact that many current environmental models are either overly abstract, or perhaps even cut across the grain of the perceptually-oriented and action-oriented environmental differentiations that are the basis of ecological task analysis. I suspect that in some cases the apparently unruly mappings that give rise to complex or probabilistic accounts of cognition can be straightened out as much by increased attention to environmental modeling as by increasingly elaborate cognitive modeling.

An Example of Ecological Task Analysis

We have performed modeling of human-environment interaction in a dynamic micro-world in order to advance approaches that could provide resources for interface design (Kirlik, Miller and Jagacinski, 1991). At a concrete level, the experimental apparatus consisted of rich, graphically displayed sources of information and a mixture of both continuous and discrete controls, similar to the kind of interface technology typical of many modern human-machine systems. At an abstract level, the task required subjects to engage in both manual control and supervisory control (Sheridan, 1987) of a set of semi-autonomous craft operating in a simulated world. The selection of a supervisory control task, in which a system operator is responsible for planning and implementing activities for (often remote) automated systems, was motivated by the introduction of automation in many existing systems and the need to design interfaces to support this form of interaction.

The experimental apparatus simulated the cockpit of a scout vehicle, over which subjects could use manual (joystick) or automatic (autopilot) control. Crews used a supervisory mode of control over four additional craft by entering strings of action commands using a text editor specifically designed for the experiment. Subjects piloted the scout within the partially forested world shown upon a dynamic, color graphical map or situation display showing the entire 100-square mile area to which activity was confined. The scout s major activity was to discover hidden objects (cargo and enemy craft) within the world. The scout was therefore equipped with a 1.5 mile radius radar for this purpose. Subjects used the additional craft primarily to act upon the discovered objects, i.e., to engage both stationary and mobile enemy craft and to load cargo and unload it at a home base. Subjects also had to attend to a number of resource management constraints (e.g., fuel, missiles, cargo capacity) in order to successfully complete each 30 minute experimental session.

The task was quite complex and many hours of practice were required to achieve mastery. However, at skilled levels of performance the selection of action was quite rapid and a fluent and often seamless mode of dynamic interaction characteristic of much skilled behavior was observed.

The *apparent* economy of behavior in this environment led to the hypothesis that subjects were relying heavily upon the rich set of graphical information as an external problem representation, with some of the attendant advantages of this processing mode as discussed above in relation to our example of the short-order cook. However, it seemed unlikely that a perception-action mode of control was possible in all cases, since some of the constraints upon productive action were not easily identifiable from the displayed information, and actions were not always available to resolve uncertainties associated with non-specific perceptual information. An ecological task analysis of this human-environment system was performed in order to identify situations where the control-display interface supported a perception-action processing mode, as well as those situations in which the interface design may have required subjects to use a more cognitively intensive mode of action selection. The results of the task analysis were used to motivate the design of a process model capable of successfully mimicking subject behavior.

An ecological task analysis of search behavior

The present example concerns modeling the selection of continuous search paths for the scout through the simulated world. A more complete description of this model as well as a description of environmental and cognitive modeling for dynamic discrete action selection in the laboratory task can be found in Kirlik, Miller, and Jagacinski (1991). Figure 4 is a depiction of a world configuration as it was displayed to subjects. The open regions, here indicated in white, were displayed in light brown. The lightly forested and heavily forested regions, here shown as light gray and dark gray, were displayed as light green and dark green, respectively. Only home base (the unfilled circle) and the initial location of the scout are shown in the figure.

Insert Figure 4 about here

Searching the world for cargo and enemy craft requires consideration of two capacities for action: scout locomotion and sighting objects with scout radar. Locomotion was most efficiently performed in open rather than forested regions due to the need to navigate around trees. Sighting objects, on the other hand, was more efficiently performed in lightly forested areas because objects were considered to be more densely located in forests. In addition, the fuel range constraints influenced search path selection since fuel expenditure rates were designed so the scout had to refuel at home base at some point during the middle third of the experimental session.

Insert Figure 5 about here

Figure 5 shows search paths created by two different subjects. Both begin by traveling north along the boundary of the inner forest, turn east to follow the top boundary of this forest, loop back west along the boundary of the upper forest, turn south along the border of the left forest, visit home base for refueling, then depart to the east and then the north, at which point the session terminated. This boundary hugging behavior resulted from the interaction of the two major criteria for search path selection: searching forests with the 1.5 mile radar to discover objects, and locomoting through open terrain to cover as much area as possible.

What is the most faithful description of the cognitive activities underlying search behavior? What is an appropriate model for the depth structure of this environment? We of course would like to find a model of depth structure that captures those environmental features to which cognition and behavior was sensitive. One could of course formulate this process as a constrained optimization problem and use a generate-and-test procedure to create alternative paths and then evaluate them with respect to an objective function. However, there are an infinite number of possible paths and the computational demands appear overwhelming. One could also attempt to describe this process in rule-based terms, using an environmental description in terms of perceptually salient objects such as forests and their borders. With this model, the cognitive processes underlying search behavior would be described in terms of the manipulation of symbols standing for discrete aspects of environmental structure.

Ecological task analysis, on the other hand, suggests that we delay assumptions about modeling environmental depth structure and cognitive processes until after the models of environmental perceptual structure and action structure have been constructed. An initial cut at constructing the perceptually-oriented environmental model required for ecological task analysis would be to describe the objects and properties perceptually apparent on this map display. However, as will be seen below, ecological task analysis suggests that we iterate and refine this initial perceptually-oriented model after constructing the action-oriented model by using the latter to help identify any initially overlooked perceptual information capable of specifying the action-oriented structure of this world. The information used as the basis for action selection may be considerably more subtle and rich than the information preserved when using perception in a purely descriptive capacity (e.g., Neisser, 1988; Bridgeman, 1991; Shebilske, 1991).

Constructing the action-oriented environmental model requires the use of action capacities as a frame of reference for environmental description. Since two action capacities underlie search behavior (locomotion and sighting objects with radar), we must describe the action-oriented structure, or affordances, of this environment for both locomotion and sighting objects. We will consider search affordances to be a simple combination of the locomotion and sighting affordances. Figure 6 shows the distribution of locomoting, sighting, and searching affordances as maps of the world in which the paths show in Figure 5 were generated. Figure 6a shows the world as it appeared on the map display. Figure 6b shows the locomoting affordance, calculated by assigning a value of zero for open regions, a value of -1.5 to lightly forested regions, and a value of -2.0 to heavily forested regions. These values were assigned to reflect the difficulty of rapidly flying the scout through these regions differing in tree density. Darker regions on the maps indicate higher affordance values.

Insert Figure 6 about here

The affordance values were selected by attempting to construct an objective measure of the degree to which relevant actions could be performed as a function of environmental structure. For example, Figure 6c shows the world sighting affordance structure. To construct this map, a four dimensional vector was associated with each world location to indicate the percentage of area that would be covered by scout radar centered at that location that was open region, lightly forested region, heavily forested region, and area beyond the world boundaries. For each point, the inner product of this vector and a sighting affordance vector was taken to determine the sighting affordance of a particular world location. The sighting affordance vector was the same for each world location and indicated the density of cargo and enemy craft within each of the four types of regions. The sighting affordance vector had a value of zero for open regions and area beyond world boundaries, and a value of 1.0 for lightly and heavily forested regions. A maximal sighting affordance would exist, therefore, in cases where the entire scout radar range covered a forested region, and a minimal sighting affordance would exist when the entire scout radar range covered either an open region or area beyond the world boundary. The graded structure of the sighing affordance distribution results from the complex interaction between the circular radar capabilities of the scout and the irregularly shaped open and forested regions that determined object density, or more generally, the interaction between the subjects action capacities and the environmental structure.

Figure 6d shows the search affordance structure, created by simply summing the values of the locomoting and sighting affordances at each world location (i.e., a location affords search if it affords both locomoting and sighting objects). This map has been rescaled to clearly indicate local optima in the search affordance structure. Considered three dimensionally, this map indicates peaks and ridges of high search affordance and valleys and holes of low searching affordance. The peak areas indicate the best compromise between the conflicting demands for locomotion through open regions and sighing objects in forested regions.

In order to define opportunities for fluent perception-actions solutions to this task, and also to define models of environmental depth structure and cognitive process, we now examine the congruence between the perceptually-oriented and action-oriented environmental models. Note that our original perceptually-oriented model differentiates the world differently than does the action-oriented model. The two models are apparently out of alignment. We now ask the question, however, what information contained within the perceptually-oriented model is available to *specify* the action-oriented structure, or affordances, in this the world? Using knowledge of human perceptual capacities together with knowledge of the displayed environmental structure, we attempt to find a way in which perception could possibly measure the displayed world in a manner that specifies the depth structure to the most faithful extent possible. What perceptual information would be necessary to specify the search affordance structure?

First, note that the perceptually-oriented model differentiates the world in an isomorphic manner to the differentiation provided by locomotion affordances. If perception can identify whether a given location is open region, light forest, or heavy forest, then information is perceptually available to fully specify locomotion affordances. To fully specify search affordances, however, perception would also have to be able to measure the sighting affordance structure. Given the manner in which sighting affordances were constructed we can define the nature of the necessary perceptual information in this case. Specifically, when foveating at a particular world location, perception would have to supply a measure of the amount of forested area within a circular area defined by the 1.5 mile radar radius of the scout. Although psychophysical experiments are surely needed to assess the degree to which this is possible (and could be straightforwardly conducted), here I will simply assume that such perceptual judgments are possible, although we may expect certain forms of systematic measurement errors. The result of such experiments would be an empirically-based perceptually-oriented model of the world that indicates the information available to specify the search affordances structure. Constructing the action-oriented environmental model, however, was necessary to first identify what kinds of psychophysical experiments to conduct.

Let us assume for the sake of this exercise that the results of such experiments suggested that people did have the perceptual ability to specify the search affordance structure. (If results indicated the subjects could not reliably estimate the search affordance structure we would then be able to define the kinds of cognitive activities necessary to do so for behavior to be productive, i.e., to be in alignment with search affordances). What would search behavior look like if subjects were simply allowing search behavior to be governed by the perceptually detected search affordances? We would perhaps expect in this case that the scout would be flown up the steepest gradient in the search affordance structure from its present position. But while the local organization in search paths may be describable in this fashion, search paths also have a global organization that is not well captured by this simple search model.

This mismatch between observed behavior and the behavior that would result from a simple perception-action solution to this task is suggestive of what kinds of additional cognitive demands this task makes upon the performer. We construe long-range or global path planning as the selection of a sequence of waypoints to be visited, where each waypoint is a peak or ridge in the search affordance structure. We still can, however, construe short-range or local navigation in terms of the simple perception-action model discussed above that results in search affordance gradient ascent. The selection of an appropriate sequence of waypoints is constrained the need to avoid backtracking through previously searched regions, and also the need to return to home base at some point in the middle third of the mission. There is no readily available perceptual information capable of specifying these constraints upon productive activity. Thus, we have a case of perceptual underspecification of action, where the subject must apparently try out a number of alternative solutions to assess the downstream effects (i.e., consistency with backtracking and fuel constraints) of each, prior to selecting a global search path.

Insert Figure 7 about here

The resulting model is shown in Figure 7. At the start of a session, the model identifies the peak areas in the search affordance map as candidate waypoints to visit during a mission. These peak areas were submitted to a generate-and-test mechanism that attempted to order the waypoints to acceptably meet backtracking and fuel constraints. Note that although a cognitively intensive process is needed for this purpose, a relatively small number of waypoints are considered, since the search affordance structure can be used to obtain a relatively low-dimensional representation of the world for long-range path planning (i.e., the set of local optima). Many of the fine details of local search affordance structure can be ignored during this process. The output of this process is an ordered sequence of waypoints. The first waypoint is then selected as a destination and was thus considered to possess an affordance for visiting. The scout did not fly in a linear path to the waypoint, however, since scout motion was determined not only by the visiting affordance but also by the local search affordance structure in the vicinity of the scout. Detailed motion commands for the scout were created by considering the perceptually-measurable search affordance structure to operate upon the scout as an attractive force field, which when combined with the a force exerted by the visiting affordance of the current waypoint, determined the direction of motion on a second by second basis. Large weights on the local search affordance values relative to the weighting on the visiting affordance provided by the current waypoint resulting in meandering motion that was very sensitive to search affordance structure. In contrast, a large weight on the visiting affordance relative to the local search affordances resulted in a direct path to the current waypoint which largely disregarded the local search affordance structure. This search model was one component in a complete process model of skilled human-environment interaction in the laboratory task. An evaluation of the empirical adequacy of the process model can be found in Kirlik, Miller, and Jagacinski (1991).

CONCLUSION

I have tried to make the case that modeling to support design requires theories which take the integrated human-environment system as the unit of analysis. Ecological task analysis is offered as a framework in which the search for the mechanisms underlying human-environment interaction can be carried out. It is quite clear that much work remains to be done in order to flesh out the details of ecological task analysis. The proposed framework is in no way a new theory of cognition, although it may be considered to be a theory of the more global phenomenon of human-environment interaction. As such, it does not contribute any new models of cognitive processes or behavior, but rather, it guides the process of modeling any particular instance of interaction, relying upon the perception-action and cognitive models the scientist may already have available. In addition, ecological task analysis provides resources for design. For an already existing design concept, it assists in the identification of cognitive demands through the process of identifying match and mismatch in perception-action environmental structure. The results of such an analysis can then specify the design of displays and controls that bring perception and action structure into better alignment.

Modeling skilled human-environment interaction that relies upon the intensive exploitation of environmental structure promises to be a messy business. We are led to this conclusion once we accept the possibility that the ability to skillfully perform within such a wide range of environments is a testimony to our adaptive abilities, rather than testimony to any abstract, context-free, general-purpose methods we may have for the selection of action. Perhaps we only resort to general-purpose methods for action selection when we fail to find (or adjust the environment to allow) more efficient perception-action solutions which are enabled by exploiting many of the concrete details of the perception and action structures of our world. The search for models of skilled interaction should not be mislead by the fact that what people can do most efficiently may be hardest to describe, and what we do most inefficiently may be easiest to describe.

Throughout this chapter I have admittedly taken a cavalier attitude with respect to a number of detailed psychological issues. This is not because I do not think such issues are important, but rather because it is all too easy (for me, at least) to get so enamored with psychological minutiae that the forest is lost for the trees. Thus, the emphasis has been on the issue how we will know a cognitive model capable of guiding environmental design when we see one. In order to take some positive steps in this direction, I have had to point out the deficiences of variety of existing modeling approaches in cognitive science. I have not done this out of lack of respect for such research. It is simply the case that the constraints that guide modeling in normal scientific practice are quite different than the constraints that must guide modeling to support design. Constraints give rise to problem decompositions, problem decompositions give rise to sets of solutions, and these solutions give rise to new problems when we aim for theoretical integration, as we must to support design. I quite expect that someone else will come along and criticize my own proposals for decomposing the cognitive modeling problem. Perhaps the social structures and physiological stressors operative in human machine systems cannot be so easily cleaved from the cognitive factors as many of us normally assume. If the failure to describe the role of such factors is a major impediment to design, we should welcome these criticisms of our theories as well.

Those with concrete design experience may claim that I have done for the design process what I have tried to guard against in the modeling process: over-rationalization of behavior. Current design practice is in many cases a far cry from the explicit problem solving search over environmental models I have portrayed. However, much of the chaos of the design process may be due to the lack of a sound theoretical basis. One must be on guard against proposing design aiding strategies that are too tightly linked to current design activity lest one run the risk of treating the symptoms rather than the cause. Obviously I believe that treating the cause requires an explicit statement of the kinds of psychological models necessary to support cognitive engineering. The design community must police its own research, for no one is likely to produce readily applicable models unless that is the explicit goal.

Acknowledgments

I am grateful to Kim Vicente for helpful comments on a previous draft of this paper. This research was supported by NASA Ames grant NAG2-656 to the Georgia Institute of Technology, Robert J. Shively, technical monitor.

References

Anderson, John R. (1990) The Adaptive Character of Thought. Hillsdale, NJ: Erlbaum.

Barwise, Jon, and Perry, John. (1983) Situations and Attitudes. Cambridge, MA: The MIT Press.

Beach, K. D. (1988) The role of external mnemonic symbols in acquiring an occupation. In M. M. Gruneberg, P. E. Morris, and R. N. Sykes, (Eds.), Practical Aspects of Memory: Current Research and Issues, Vol. I. New York: Wiley.

Birmingham, H. P. and Taylor, F. V. (1954) A design philosophy for man-machine control systems. Proceedings of the I.R.E., Vol. XLII, No. 12, (pp. 1748-1758).

Bohlman, L. (1979). Aircraft accidents and the theory of the situation. Resource management the flight deck. Proceedings of a NASA/industry workshop. NASA Conference Proceedings 2120.

Braitenberg, Valentino. (1984) Vehicles: Experiments in Synthetic Psychology. Cambridge, MA: The MIT Press.

Brehmer, B., and Joyce, C. R. B., Eds., (1988) Human Judgment: The SJT View. New York: North Holland.

Bridgeman, Bruce (1991) Separate visual representations for perception and for visually guided behavior. In Pictorial Communication in Virtual and Real Environments, Stephen R. Ellis, Mary K. Kaiser, and Arthur J. Grunwald, (Eds.). New York: Taylor & Francis, (pp. 317-327).

Brunswik, Egon. (1952) The conceptual framework of psychology. In, International Encyclopedia of Unified Science (Vol. 1, No. 10). Chicago, IL: University of Chicago Press.

Brunswik, Egon. (1956) Perception and the Representative Design of Psychological Experiments. University of California Press.

Card, Stuart K., and Newell, Allen. (1989) Cognitive architectures. In Jerome I. Elkind, Stuart K. Card, Julian Hochberg, and Beverly Messick Huey, (Eds.), Human Performance Models for Computer-Aided Engineering. Washington, D.C.: National Academy Press, (pp. 173-179).

Carroll, J. M., and Campbell, R. L. (1989) Artifacts as psychological theories: The case of human-computer interaction. Behaviour and Information Technology, 8, (pp. 247-256).

Carroll, John M. (1991) Introduction: The Kittle House manifesto. In Designing Interaction, John M. Carroll, (Ed.). New York: Cambridge University Press, (pp. 1-16).

Claxton, Guy. (1988) How do you tell a good cognitive theory when you see one? In G. Claxton, (Ed.), Growth Points in Cognition. New York: Routledge.

Donald, Merlin. (1991) Origins of the Modern Mind. Cambridge, MA: Harvard University Press.

Edwards, E. (1988). Introductory overview. In E.L. Wiener and D.C. Nagel (Eds.) Human Factors in Aviation. New York: Academic Press.

Fischoff, B., Slovic, P., and Lichtenstein, S. (1978) Fault trees: Sensitivity of estimated failure probabilities to problem representation. Journal of Experimental Psychology: Human Perception and Performance, 4, (pp. 330-344).

Flach, John M. (1990a) Control with an eye for perception: Precursors to an active psychophysics. Ecological Psychology, Vol. 2, No. 2, (pp. 83-112).

Flach, John M. (1990b) The ecology of human-machine systems I: Introduction. Ecological Psychology, 2(3), (pp. 191-205).

Gibson, James J. (1979) The Ecological Approach to Visual Perception. Boston, MA: Houghton-Mifflin.

Gibson, James J. (1966) The Senses Considered as Perceptual Systems. Boston, MA: Houghton-Mifflin.

Gibson, James J. (1967/1982) James J. Gibson autobiography. In E. Reed and R. Jones, Reasons for Realism. Hillsdale, NJ: Erlbaum, (pp. 7-22).

Hammond, Kenneth R., Hamm, Robert M., and Grassia, Janet. (1986) Generalizing over conditions by combining the multitrait-multimethod matrix and the representitive design of experiments. Psychological Bulletin, Vol. 100, No. 2, (pp. 257-269).

Hammond, Kenneth R., Stewart, Thomas R., Brehmer, Berndt, and Steinmann, Derick O. (1975) Social judgment theory. Center for Research on Judgment and Policy Report No. 176, University of Colorado Institute for Behavioral Science.

Hammond, Kenneth R., Hamm, Robert M., Grassia, Janet, and Pearson, T. (1987) Direct comparison of analytical and intuitive cognition in expert judgment. IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-17, No. 5, (pp. 753-770).

Hammond, Kenneth R. (1955) Probabilistic functioning and the clinical method. Psychological Review, 62, (pp. 255-262).

Hutchins, Edwin. (1991) How a cockpit remembers its speeds. Paper presented at the Second NASA Ames Aviation Safety/Automation Program Investigators Meeting eeting. August 14-16, 1991.

Hutchins, Edwin. (1988) The technology of team navigation. Institute for Cognitive Science Report 8804, University of California, San Diego, La Jolla, CA.

Johnson-Laird, P. N. (1975) Models of deduction. In R. J. Falmagne, (Ed.), Reasoning: Representation and Process in Children and Adults. Hillsdale, NJ: Erlbaum.

Kahneman, D. and Tversky, A. (1979) Prospect theory: An analysis of decisions under risk. Econometrica, 47, (pp. 263-291).

Kirlik, Alex, Miller, R. A., and Jagacinski, Richard J. (1991) Supervisory control in a dynamic uncertain environment II: A process model of skilled human-environment interaction. Manuscript submitted for publication.

Klein, Gary A. (1989) Recognition-primed decisions. In W. B. Rouse, (Ed.), Advances in Man-Machine Systems Research. Greenwich, CT: JAI Press, (pp. 47-92).

Kotovsky, K., Hayes, J. R., and Simon, H. A. (1985) Why are some problems hard? Evidence from the tower of Hanoi. Cognitive Psychology, 17.

Larkin, Jill H., and Simon, Herbert A. (1987) Why a diagram is (sometimes) worth a thousand words. Cognitive Science, Vol. 11, (pp. 65-99).

Marr, David. (1982) Vision. New York: Freeman.

McRuer, D. T. and Jex, H. R. (1967) A review of quasi-linear pilot models. IEEE Transactions on Human Factors in Electronics, HFE-8, (3), (pp. 231-249).

Meister, David. (1989) Conceptual Aspects of Human Factors. Baltimore, MD: The Johns Hopkins University Press.

Meister, David. (1984) Letter. Human Factors Society Bulletin, 27(10), (p. 2).

Neisser, Ulric. (1989) Direct perception and recognition as distinct perceptual systems. Address presented to the Cognitive Science Society, August, 1989.

Neisser, Ulric. (1976) Cognition and Reality. New York: W.H. Freeman and Company.

Newell, Allen, and Simon, Herbert A. (1972) Human Problem Solving. Englewood Cliffs, NJ: Prentice-Hall.

Norman, Donald A. (1986) Cognitive engineering. In D. A. Norman and S. W. Draper, (Eds.), User Centered System Design. Hillsdale, NJ: Erlbaum.

Norman, Donald A. (in preparation) Things That Make Us Smart. Reading, MA: Addison-Wesley.

Rasmussen, Jens. (1986) Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering. New York: North-Holland.

Reitman, Walter, Nado, Robert, and Wilcox, Bruce. (1978) Machine perception: What makes it so hard for computers to see? In C. Wade Savage (Ed.), Perception and Cognition Issues in the Foundations of Psychology, Minnesota Studies in the Philosophy of Science, Volume IX. Minneapolis, MN: University of Minnesota Press, (pp. 65-88).

Rouse, William B. (1987) Much ado about data. Human Factors Society Bulletin, 30(9), (pp. 1-3).

Sarter, Nadine B. and Woods, David D. (1991) Situation awareness: A critical but ill-defined phenomenon. The International Journal of Aviation Psychology, 1, (1), (pp. 45-57).

Shebilske, Wayne, L. (1991) Visuomotor modularity, ontogeny and training high-performance skills with spatial instruments. In Pictorial Communication in Virtual and Real Environments, Stephen R. Ellis, Mary K. Kaiser, and Arthur J. Grunwald, (Eds.). New York: Taylor & Francis, (pp. 305-315).

Sheridan, Thomas B. (1987) Supervisory control. In G. Salvendy, (Ed.), Handbook of Human Factors. New York: John Wiley and Sons, (pp. 1243-1268).

Suchman, Lucy. (1987) Plans and Situated Actions: The Problem of Human-Machine Communication. New York: Cambridge University Press.

Tversky, A. and Kahneman, D. (1981) The framing of decisions and the psychology of choice. Science, 211, (pp. 453-458).

Wason, P. C. and Johnson-Laird, P. N. (1972) Psychology of Reasoning: Structure and Content. Cambridge, MA: Harvard University Press.

Whiteside, J., Bennett, J., and Holtzblatt, K. (1988) Usability engineering: Our experience and evolution. In M. Helander, (Ed.), Handbook of Human-Computer Interaction. New York: North Holland, (pp. 791-817).

Wise, James A. (1985) Decisions in design: Analyzing and aiding the art of synthesis. In Behavioral Decision Making, George Wright (Ed). New York: Plenum, (pp. 283-308).

Woods, David D. and Roth, E. M. (1988) Cognitive systems engineering. In M. Helander (Ed.), Handbook of Human-Computer Interaction. New York: North Holland, (pp. 3-43).

Figure Captions

Figure 1. The Lens model of Brunswik.

Figure 2. The Ecological task analysis framework. Like the Lens model framework for human judgment, the ecological task analysis framework represents the integrated human-environment system in a symmetrical arrangement. Matches between the perceptually-oriented and action-oriented environmental differentiations indicate opportunities for the fluent perceptual guidance of activity. Structural mismatches between the models of surface perceptual structure and surface action structure are indicative of demands for cognitive activity, and corresponding models of the environmental depth structure.

Figure 3. Four possible results of an ecological task analysis. Each case reflects a different type of match or mismatch in environmental perception and action strucuture, and is thus indicative of a different form of cognitive demands, and a different type of remedial interface design solution.

Figure 4. Map display showing home base, the scout, and forested areas.

Figure 5. Search paths produced by two different subjects in the same world configuration.

Figure 6. Four representations of the same world. (a) Displayed representation. (b) Locomotion affordance representation. (c) Sighting affordance representation. (d) Search affordance representation scaled to emphasize local optima in the affordance structure.

Figure 7. Process model of search path generation. Local optima in the search affordance structure are organized into a series of waypoints via a generate-and-test scheduler. Detailed scout motion is produced by a combination of attractive forces from both the current waypoint and the local search affordance structure in the vicinity of the scout.