Ecological Resources for Modeling Interactive Behavior and Embedded Cognition

Alex Kirlik

Human Factors Division and Beckman Institute

University of Illinois at Urbana-Champaign

In (2007) W. Gray (Ed.), *Integrated Models of Cognitive Systems*. New York: Oxford University Press.

Ecological Resources for Modeling Interactive Behavior and Embedded Cognition

Alex Kirlik

ABSTRACT

A recent trend in cognitive modeling is to couple cognitive architectures with computer models or simulations of dynamic environments to study interactive behavior and embedded cognition. Progress in this area is made difficult by the fact that cognitive architectures traditionally have been motivated by data from discrete experimental trials using static, noninteractive tasks. As a result, additional theoretical problems must be addressed to bring cognitive architectures to bear on the study of cognition in dynamic and interactive environments. I identify and discuss three such problems dealing with the need to model the sensitivity of behavior to environmental constraints, the need to model context-specific adaptations underlying expertise, and the need for environmental modeling at a functional level. I illustrate these problems and describe how we have addressed them in our research on modeling interactive behavior and embedded cognition.

INTRODUCTION

An emerging trend in the study of interactive behavior and embedded cognition is to couple a cognitive model implemented in a cognitive architecture with a computational model or simulation of a dynamic and interactive environment such as a flight simulator, military system or videogame (Gluck, this volume; Gluck & Pew, 2005; Foyle & Hooey, in press; Gray, Schoelles & Fu, 2000; Shah, Rajyagura, St. Amant and Ritter, 2003; Salvucci, this volume). Some of the impetus for this research is a growing interest in prospects for using computational cognitive modeling as a technique for engineering analysis and design. These attempts follow by about a generation a set of related attempts to model closed-loop cognition and behavior in the field of human-machine systems engineering (Rouse, 1984; 1985; Sheridan & Johannsen, 1976; also see Pew, this volume). As noted by Sheridan (2002), these human-machine systems engineering models represented a desire "to look at information, control, and decision making as a continuous process within a closed loop that also included physical subsystems – more than just sets of independent stimulus-response relations" (Sheridan, 2002, p. 4).

The cognitive architectures available to today's modeling community such as ACT-R (Anderson, this volume), COGENT (Cooper, this volume), ADAPT (Doane, this volume), EPIC (Hornoff, this volume), Soar (Ritter, this volume), or Clarion (Sun, this volume) are better suited than were their engineering-based predecessors for describing the internal processes underlying behavior beyond merely "sets of independent of stimulus-response relations." (Sheridan, 2002, p. 4). So why is it still so difficult to model a (typically experienced) pilot, driver or videogame player with a cognitive architecture? My aim in this chapter is to address this question by providing some distinctions and modeling techniques that will hopefully accelerate progress in modeling interactive behavior and embedded cognition.

THEORETICAL ISSUES IN MODELING EMBEDDED COGNITION

Difficulties in what is sometimes called “scaling up” cognitive modeling to the complexities of dynamic and interactive contexts such as aviation and driving largely have their origins in tasks and data. In particular, there are qualitative differences between the types of tasks and data sets that gave rise to many of the better known cognitive architectures and the types of tasks and data sets characteristic of many dynamic and interactive contexts. A central goal of this chapter is to bring some clarity to the description of these qualitative differences and their implications. My hope is that clarifying these distinctions will be useful in moving beyond vague and not particularly informative discussions on the need to “scale up” modeling, to bridge theory and application, or even worse, to move from the laboratory to the “real” world.

As I will try to show in the following, what is at issue here is not so much a scaling up as a scaling over. Modeling interactive behavior and embedded cognition raises interesting and challenging theoretical questions that are distinct from the types of theoretical questions that provided the traditional empirical foundation for cognitive architectures. By “distinct” I mean that many of the theoretical questions that arise when modeling dynamic and interactive tasks are not reducible in any interesting sense to the questions that motivated the design of many current cognitive architectures. New and different questions arise, along with their attendant modeling challenges and opportunities.

In the following sections I discuss three types of theoretical issues that emerge when examining mismatches between the types of empirical data that have typically motivated the design of cognitive architectures and the types of data confronting modelers of interactive and embedded cognition in operational contexts. The first issue deals with the fact that cognitive architectures have chiefly been designed to model cognition in discrete and static tasks (i.e., laboratory trials) whereas data on embedded cognition often reflects performance in continuous and dynamic tasks. I suggest that modeling cognition and behavior in the latter type of tasks creates a need to model the manner in which behavior is dynamically sensitive to environmental constraints and opportunities. Doing so may require expanding one's view of the functional contribution of perception to intelligent behavior. Rather than viewing perception to be devoted solely to reporting the existence of objects and their properties to cognition in objective or task-neutral terms, it may be increasingly important to also view perception as capable of detecting information that specifies opportunities for behavior itself.

The second issue concerns the fact that the design of cognitive architectures has mainly been motivated by data from largely task-naive subjects, or often with subjects with no more than a few hours of task-relevant experience. In contrast, modeling cognition in operational contexts such as aviation and driving often involves data from highly experienced performers. It is impossible to create a good model of a performer who knows more about the task environment than does the modeler. As a result, modeling experienced cognition requires not only expertise in cognitive modeling but also an ability to obtain expert knowledge of the relevant task and environment. While modeling students acquiring Lisp programming or arithmetic skills allows one to obtain this expert knowledge from books, modeling performers in interactive and dynamic domains typically requires detailed empirical study (e.g., Gray & Kirschenbaum, 2000). This knowledge is required not only to guide the development of a cognitive model, but also to develop a formal model of the task environment[1] with which the cognitive model can interact.

Finally I discuss theoretical questions that arise out of the profoundly interactive nature of

[1] At the outset I wish to stress that my many comments in this chapter on the need for environmental modeling refer to models of the external world of the performer and not to models of the performer's internal representations of that world, although the latter may also be needed.

much behavior and embedded cognition in operational contexts. In particular I suggest that interactive tasks create a need to view and model the environment much more functionally than may be required when modeling noninteractive contexts. This suggests that a largely physicalistic approach to environmental modeling, for example in terms of the types, locations and features of perceptible objects on a display is likely to be insufficient for understanding cognition and behavior as a functional interaction with the world. Richer techniques for functional-level environmental modeling are needed to marry the functional accounts of cognition provided by cognitive modeling with functional accounts of the environment. When modeling interactive behavior and embedded cognition, one can get only so far by trying to couple functional models of cognition with physical models of the environment. A functional perspective must be adopted for both.

After each of the three issues outlined above is discussed in greater detail I then present a set of modeling projects from our previous research touching in one way or another on these issues. Each of these projects represents an explicit attempt to computationally model interactive behavior and embedded cognition in a dynamic and interactive environment.

Modeling Sensitivity to Environmental Constraints and Opportunities

One axiom within the engineering-oriented modeling tradition discussed previously concerned the necessity of modeling the environment as a prerequisite to modeling cognition and behavior. As Baron (1984) put it:

> Human behavior, either cognitive or psychomotor, is too diverse to model unless it is sufficiently constrained by the situation or environment; however, when these environmental constraints exist, to model behavior adequately, one must include a model for that environment. (Baron, 1984, p.6)

Baron's comment places a spotlight on the *constraining* (note: not *controlling*) nature of the environment as an important source of variance that must be known when modeling behavior. Understanding how environmental constraints and opportunities determine the playing field of behavior is such a mundane exercise in everyday life that we often forget or overlook the important role that it plays. You will obviously not be swimming in the next minute unless you are already sitting near a pool or on a beach. In experimental research, a modeler typically would not get any credit for explaining all the variance associated with things that our subjects do and not do because a task does and does not provide the opportunity to do those things. Instead, the focus is on explaining variance above and beyond what could be "trivially" predicted by examining the carefully equated opportunities for behavior an experiment affords.

All of the cognitive architectures of which I am aware, due to their origins in describing data from experimental psychology, have built into them this focus on explaining variance in behavior above and beyond environmental constraints on that behavior. This can be seen from what these models predict: reaction times which, if the experiment is "well designed," represent solely internal constraints but not external task constraints (a potential confound); the selection of an action from a set of actions all of which are carefully designed to be equally available to the subject (another potential confound). Cognitive experimentalists typically take great pains to equate the availability of the various actions (e.g., keypresses) presented to participants. It is easy to overlook how this tenet of experimental design limits generalization to contexts in which the detection of action opportunities themselves and variance associated with the possibly differing levels of the availability of various actions contribute to variance in behavior.

I am hardly the first to note the many differences between the largely static, noninteractive environment of the discrete laboratory trial and environments such as videogames, aviation and driving. But note the implications regarding the necessity of environmental modeling in the two cases. To explain variance in the static laboratory experiment, since credit is given only for explaining or predicting variance above and beyond what is environmentally constrained, no attention need be given to modeling how behavioral variance is environmentally constrained. As such, cognitive architectures typically provide no resources explicitly dedicated to this ubiquitous aspect of cognition and behavior in everyday situations. In modeling experimental data, determining which actions are appropriate given the environmental context is a task performed *by the modeler* and encoded once and for all in the model: it is rarely if ever a modeled inference. This only works because the environment of the laboratory trial is presumed to be static in the sense that all (relevant) actions are always equally available.

So the modeler who would like to apply cognitive architectures motivated almost solely by data from such experiments to dynamic, interactive situations is largely on his or her own when determining how to make the model sensitive to environmental constraints and opportunities in a dynamic and interactive fashion. Modeling this type of sensitivity will be necessary any time a performer is interacting with a dynamic, and especially uncertain environment. Both dynamism and uncertainty place a premium on perception to aid in determining the state of the environment in terms of which behaviors are and are not appropriate at a given time. As such, the modeler will be faced with questions concerning the design of perceptual mechanisms to aid in performing this task (e.g., Fajen & Turvey, 2003). If “primitive” perceptual mechanisms are provided by the architecture, the modeler will be faced with questions about which environmental information these mechanisms should be attuned to, and additional “primitive” mechanisms may need to be invented (e.g., Runeson, 1977). This may well require reference to an environmental model that represents perceptually available information at a high level of fidelity, and the task of defining perceptual units or objects may present nontrivial problems. All of these issues speak to the question of why it has proven to be difficult to use computational cognitive architectures to model performers in dynamic, interactive environments.

Knowing as Much or More than the Performer

I have already discussed perhaps the most primitive aspect of adaptation to an environment: ensuring that behavior is consistent with environmental constraints on behavior. Assume for a moment that this problem is solved and we are interested solely in examinations of cognition and behavior above and beyond what is so constrained. One finding from the human-machine systems tradition discussed previously is that a good step toward predicting the behavior of experienced performers in dynamic, interactive contexts is to analyze a task in terms of what behavior would be optimal or most adaptive (see Pew, this volume). At first blush this approach would seem to dovetail quite nicely with modeling approaches with origins in either rational analysis (Anderson, this volume) or ecological rationality (Todd, this volume).

It is important to note, however, that appeals are made to different quarters when one assumes the rationality or optimality of basic cognitive mechanisms and when one assumes the rationality or optimality of experienced behavior. The rationality underlying the design of ACT-R's memory, categorization and inference mechanisms and Gigerenzer and Todd's (1999) toolbox of fast and frugal heuristics appeals to evolutionary arguments rather than to learning or experience per se. The subjects in experiments performed from the perspective of both these adaptive approaches to cognition are not typically presumed to have any first hand experience with the tasks studied. The hypothesis that memory exhibits a Bayesian design or that some decisions are made by a recognition heuristic are intended as claims about the human cognitive architecture independent of any *task-specific* experience. In fact, one can look at learning to be accumulating the additional adaptations necessary to perform a given task like an experienced performer instead of like a task-naive novice.

Much, if not most, modeling research done in dynamic, interactive environments is oriented toward understanding and supporting skilled performance. Much, if not most, experimental research done to inform the design of cognitive architectures uses largely task-naive subjects, or at best subjects with only a few hours of instruction or training. It is hardly surprising, then, that researchers interested in modeling the behavior of automobile drivers, videogame players and pilots have to invent their own methods for identifying and codifying the experiential adaptations underlying skilled behavior. This is true even if they select and use a cognitive architecture informed by rationality or optimality considerations, and even if the behavior to be modeled is highly rational or even optimal.

Modeling task-naive behavior can be done by similarly task-naive scientists. The main requirement is expertise in cognitive modeling. But modeling expert performance also requires expert knowledge of the task environment to which the expert is adapted. Neisser (1976) put the matter of modeling expert performance as follows:

> What would we have to know to predict how a chess master will move his pieces, or his eyes? His moves are based on information he has picked up from the board, so they can only be predicted by someone who has access to the same information. In other words, an aspiring predictor would have to understand the position at least as well as the master does; he would have to be a chessmaster himself! If I play chess against the master he will always win, because he can predict and control my behavior while I cannot do the reverse. To change this situation I must improve my knowledge of chess, not of psychology. (Neisser, 1976, p. 183)

Our own experiences in modeling experienced performers, detailed in the examples to follow, has taught us that one must often spend as much, if not more, time studying and formally modeling the external task environment than is spent modeling inner cognition. One reason this is required to enable modeling of the highly context-specific cognitive adaptations underlying expertise.

Mind and World Function in Concert

In his wonderfully researched and written biography of the late Nobel Prize winning physicist Richard Feynman, James Gleick relates an episode in which MIT historian Charles Weiner was conducting interviews with Feynman at a time when Feynman had considered working with Weiner on a biography. Gleick writes that Feynman, after winning the Nobel Prize, had begun dating his scientific notes, "something he had never done before" (Gleick,1992, p. 409). In one discussion with Feynman, "Weiner remarked casually that his new parton notes represented 'a record of the day-to-day work,' and Feynman reacted sharply" (ibid, p. 409). What was it about Weiner's comment that drew a "sharp" reaction from this great scientist? Did he not like his highly theoretical research described merely as "day-to-day work"?

No, and the answer to this question reflects, to me at least, something of Feynman's ability to have deep insights, not only into physics, but into other systems as well. Feynman's reaction to Weiner describing his notes as "a record" was to say: "I actually did the work on the paper." (ibid, p. 409). To which an apparently uncomprehending Weiner responded, "Well, the work was done in your head, but the record of it is still here." (ibid, p. 409). One cannot fail to sense frustration in Feynman's retort: "No, it's not a *record*, not really. It's *working*. You have to work on paper, and this is the paper. Okay?" (ibid, p. 409, italics in the original).

My take on this interchange is that Feynman had a deep understanding of how his work was composed of a functional transaction (Dewey, 1896) between his huge accumulation of internal cognitive tools as well as his external, cognitive tools of pencil and paper, enabling him to perform functions such as writing, reflecting upon, and amending equations, diagrams, and so on (cf. Donald, 1991, Vygotsky, 1981). Most importantly, note Feynman's translation from Weiner's description of the world in terms of physical form ("No, it's not a *record*, not really") into a description in terms of function ("It's working").

Why did Weiner have such a difficult time understanding Feynman? External objects, such as Feynman's notes, do of course exist as things, typically described by nouns. Yet, in our functional transactions with these objects, the manner in which they contribute to cognition and behavior requires that these things also be understood in functional terms, that is, in terms of their participation in the operation of the closed-loop, human-environment system (cf. Monk, 1998, on "cyclic interaction"). Weiner, like so many engineering students through the ages, apparently had difficulty in viewing the external world not only in terms of form (nouns) but also in terms of function (verbs).

I share this anecdote here because I believe it to be an exceptional illustration of the fact that studying expert behavior not only presents challenges for understanding what the expert knows, but also challenges for understanding how the expert's environment contributes to cognition and how that contribution should be described (Hutchins, 1995). As the examples presented below will demonstrate, we have found in our own modeling of interactive behavior and embedded cognition a need to understand a performer's environment in functional terms, as a dynamic system in operation. Human-environment interaction is then understood in terms of a functional coupling between cognition and the environment functionally described. When modeling experienced performers engaged in interactive behavior and embedded cognition, I suggest that one has a much greater chance of identifying regularities in behavior by analysis at the functional level than by searching for these regularities in patterns of responses to stimuli described in physical terms. Modeling the environment in functional terms is also critically important when trying to model how a person might use tools in the performance of cognitive tasks, as the following examples will hopefully demonstrate.

I highlight the importance of adopting a functional perspective on environmental modeling for a number of reasons. As mentioned in the opening of this chapter, a trend currently exists to couple models with simulations of dynamic and interactive environments such as flight simulators, videogames and the like. While this is an important technical step in the evolution of cognitive modeling, having such an external simulation of course does not obviate the need for addressing the theoretical problem of modeling the environment in functional terms relevant to psychology. A bitmap model of the visual environment, for example, could be helpful in identifying the information available to a model's perceptual (input) mechanisms. This environmental model, however, is insufficient for determining what information a model should be perceiving at what time in order to mimic human cognition and performance.

I have little else to say about the importance of functional modeling of the environment at a general level other than to alert the reader to attend to its prevalence in the modeling examples that follow. These examples hopefully demonstrate how functional analysis allowed us to gain at least some insight into issues such as:

- Timing issues associated with the dynamic coupling of cognition and environment;
- How skilled performers might come to perceive an environment in functional terms i.e. as opportunities for action;
- How complex behavior can arise from the coupling of simple heuristics with a complex

environment;

- How people might functionally structure their environment to reduce cognitive demands;
- How making cost-benefit analyses of decision making may require extremely task-specific computations of environmental contingencies;
- How human error might arise from generally adaptive heuristics operating in ecologically atypical situations.

Models illustrating these points and others are described in the following section.

MODELING INTERACTIVE BEHAVIOR AND EMBEDDED COGNITION

In this section I describe a set of cognitive models sharing a few common themes. Each represents an attempt to computationally model human cognition and behavior in a dynamic and interactive environment. None of the models were created in an attempt to develop a unified cognitive architecture. Instead, the central reason modeling was performed was to try to shed light on how experienced performers could have possibly managed to meet the demands of what we believed to be extremely complex dynamic and interactive tasks. In other words, in none of these cases were we in the possession of knowledge of how the task could even possibly be performed in a manner consistent with known cognitive limitations prior to analysis and modeling.

Our focus on modeling *experienced* performers in *dynamic* and *interactive* tasks placed a premium on addressing the three theoretical questions discussed earlier in this chapter concerning sensitivity of behavior to environmental constraints, the need to identify and describe experiential adaptations and the need for detailed functional analysis and modeling.

The Scout World: Modeling the Environment with Dynamic Affordance Distributions

The first modeling example illustrates the use of a finely grained, functional description of an environment in terms of Gibson's (1979) theory of affordances; i.e., a functional description of the environment in terms of opportunities for action. This study shed light into understanding the fluency of behavior in a highly complex, dynamic task, plausible explanations of the differences between high and low performers, and insights into why we believe that some knowledge underlying skill or expertise may appear to take on a tacit (Polanyi, 1966), or otherwise unverbalizable, form.

Consider Figure 1, which depicts an experimental participant performing a dynamic, interactive simulation of a supervisory control task described here as the *Scout World*. This laboratory simulation required the participant to control not only his or her own craft, called the Scout, but also four additional craft over which the participant exercised supervisory control (Sheridan, 2002), by entering action plans at a keyboard (e.g., fly to a specified waypoint,

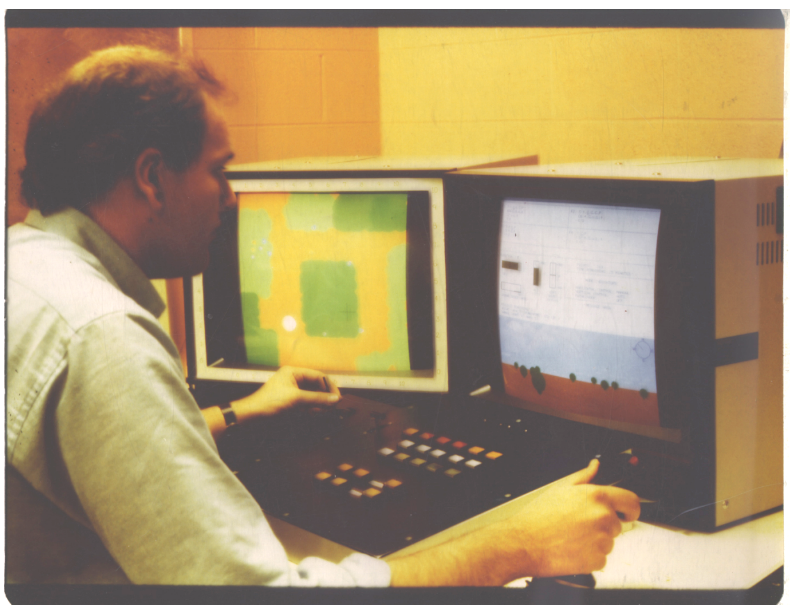

Figure 1. Experimental Participant Performing the *Scout World* Task

conduct patrol, load cargo, return to a home base, etc.). The left monitor in Figure 1 depicts a top-down situation display of the partially forested, 100-square mile world to which activity was confined. The display on the right shows an out-the-window scene (lower half) and a set of resource and plan information for all vehicles under control (upper half). The participant's task was to control the activities of both the Scout and the four other craft to score points in each 30-minute session by processing valued objects that appeared on the display once sighted by Scout radar. See Kirlik, Miller and Jagacinski (1993) for details.

Our goal was to create a computer simulation capable of performing this challenging task, and one that would allow us to reproduce, and thus possibly explain, differences between the performance of both one- and two-person crews, and novice and expert crews. At the time, the predominant cognitive modeling architectures, such as Soar (Ritter, this volume, Newell, 1990), ACT-R (Anderson, this volume; Anderson & Lebiere, 1998), and the like did not have mature perception and action resources allowing them to be coupled with external environments, nor had they been demonstrated to be capable of performing dynamic, uncertain, and interactive tasks (a limitation Newell agreed to be a legitimate weakness of these approaches: see Newell, 1992). In addition, modeling techniques drawn from the decision sciences would have provided an untenably enumerative account of participants' decision processes, and were rejected due to bounded rationality considerations (Simon, 1956).

Instead, and what was a relatively novel idea at the time, we observed that our participants seemed to be relying heavily on the external world (the interface) as "its own best model" (Brooks, 1991). This was suggested not only by intimate perceptual engagement with the displays, but also by self-reports (by participants) of a challenging, yet deeply engaged and often enjoyable sense of "flow" (Csikszentmihalyi, 1993) during each 30-minute session (not unlike any other "addictive" videogame or sport). We thus began to entertain the idea that if we were going to model the function of our human performers, we would have to model their world in functional terms as well, if we were to demonstrate how the two functioned collectively, and in concert. This turned us to the work of Gibson (1979), whose theory of affordances provided an account of how people might be attuned to perceiving the world functionally; in this case, in terms of actions that could be performed in particular situations in the *Scout World.*

Following through on this idea entailed creating descriptions of the environment using the experimental participant's capacities for action as a frame of reference to achieve a functional description of the *Scout World* environment. That is, instead of creating solely perceptually-oriented descriptions in terms of, say, object locations and colors, we described spatiotemporal regions or "slices" of the environment as "fly-throughable," "land-onable," "load-able" and so on. A now classic example of this technique was presented by Warren (1984), who measured the riser heights of various stairs in relation to the leg lengths of various stair climbers and found, in this ratio, a functional invariance in people's ability to perceptually detect whether a set of stairs would be climbable (for them) or not. Warren interpreted this finding to mean that people could literally perceive the "climbability" of the stairs; i.e., that people can perceive the world, not only in terms of form, but also in terms of function.

Like Warren, we created detailed, quantitative models of the *Scout World* environment in terms of the degree to which various environmental regions and objects afforded locomotion, searching (discovering valued objects by radar), processing those objects (loading cargo, engaging enemy craft), and returning home to unload cargo and reprovision. Because participants' actions influenced the course of events experienced, they shaped or partially determined the affordances of their own worlds. Flying the scout through virgin forest to sight and discover cargo, for example, created new action opportunities (cargo loading), and once cargo were loaded these opportunities in turn ceased to exist. In such situations the state of the task environment is, in experimental psychology terminology, both a dependent and independent variable[2]. This observation is useful in coming to understand the need for functional-level modeling of the environment to describe closed-loop dynamics. Not only must "S-R" relations be described (with some theory of cognition), but so must "R-S" relations be described, the latter requiring a model of environmental dynamics to depict how the environment changes as a function of human activity.

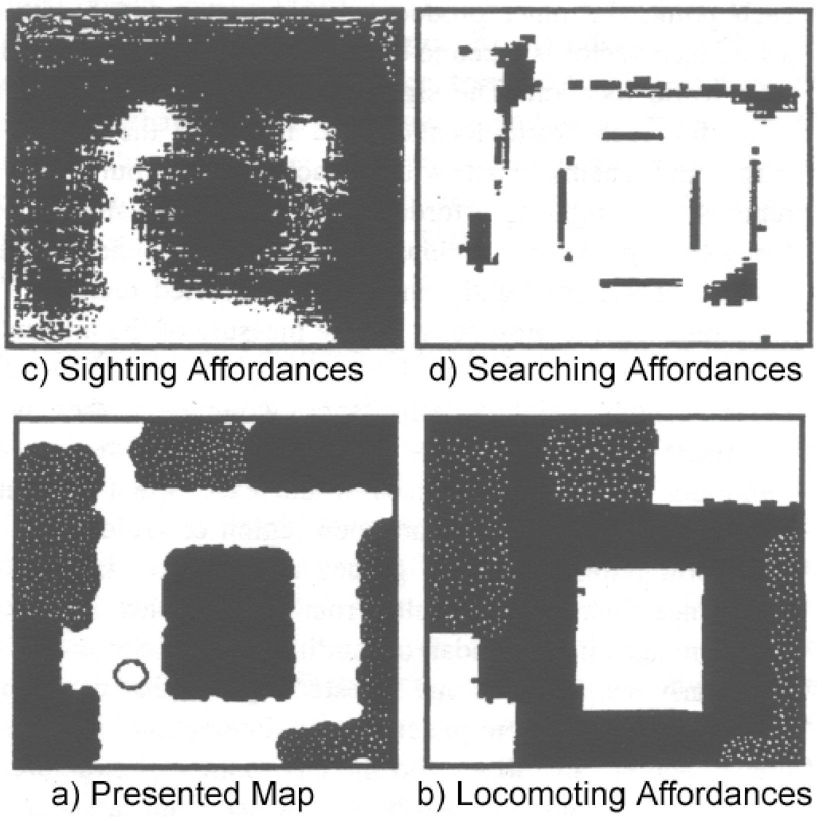

Figure 2. The Presented World Map (a), a Map of Affordances for Locomotion (b), a Map of Affordances for Sighting Objects (c), and a Final Searching Affordance Map (d)

Figure 2 contains a set of four maps of the same *Scout World* layout, including a representation purely in terms of visual form, and as shown to participants (a), and functional representations in terms of affordances for actions of various types (b, c, d). For the Scout, for

[2] Actually, I believe this raises the question of whether the logic underlying the notion of "independent" and "dependent" variables is even appropriate in such situations (Dewey, 1896) but this is beyond the current scope.

example, locomotion (flying) was most readily afforded in open, unforested areas (the white areas in Figure 2a), and less readily afforded as forest density grew. As such, Figure 2b shows higher locomotion affordances as dark and lower affordances as lighter. (Here we are of course simply using grayscale coding to represent these affordance values to the reader; in the actual model, the “dark” regions had high quantitative affordance values, and the “light” regions had relatively low quantitative affordance values.) Since the Scout radar for sighting objects (another action) had a 1.5 mile radius, and valued objects were more densely scattered in forests, the interaction between the Scout’s capacity for sighting and the forest structure was more graded and complex, as shown in Figure 2c (darker areas again indicating higher sighting affordance values). Considering that the overall affordance for searching for objects was comprised of both locomotion and sighting affordances (searching was most readily afforded where one can most efficiently locomote and sight objects), the final searching affordance map in Figure 2d was created by superimposing Figures 2b and 2c. Figure 2d thus depicts ridges and peaks that maximally afforded the action of searching.

As explained in Kirlik et al. (1993), this functional, affordance-based differentiation of the environment provided an extremely efficient method for mimicking the search paths created by participants. We treated the highest peaks and ridges in this map as successive waypoints that the Scout should attempt to visit at some point during the mission, thus possessing an attractive “force.” Detailed scout motion was then determined by a combination of these waypoint forces and the entire, finely graded, search affordance structure, or field. As one might expect, placing a heavy weight on the attractive forces provided by the waypoint peaks (as opposed to the entire field of affordances) resulted in scout motion that looked very goal-oriented in its ignorance of the immediately local search affordance field. On the other hand, reversing these weights resulted in relatively meandering, highly opportunistic scout motion that was strongly shaped by the local details of the finely grained search affordance field.

In an everyday situation such as cleaning one's house, the first case would correspond to rigidly following a plan to clean rooms in a particular order, ignoring items that could be opportunistically straightened up or cleaned along the way. The second case would correspond to having a general plan, but being strongly influenced by local opportunities for cleaning or straightening up as one moved through one's house. In the actual, computational *Scout World* model, this biasing parameter was set in a way that resulted in scout search paths that best mimicked the degree of goal-directedness versus opportunism in the search paths observed.

For object-directed rather than region-directed actions, such as loading cargo or visiting home base, the *Scout World's* affordances were centered on those objects rather than distributed continually in space. As shown in Figure 3, we created a set of dynamic affordance distributions for these discrete, object-directed actions for both the Scout and the four craft under supervisory control (F1 – F4 in Figure 5a). Each of the 15 distributions shown in Figure 3a indicates the degree to which actions directed toward each of the environmental objects that can be seen in Figure 3b were afforded at a given point in an action-based (rather than time-based) planning horizon. Space precludes a detailed explanation of how these distributions were determined (see Kirlik et al., 1993 for more detail). To take one example, consider the craft "F1" over which the participant had supervisory control by entering action plans via a keyboard. F1 appears in the northwest region of the world as shown in Figure 5b, nearby a piece of cargo labeled "C1". The "First Action" affordance distribution for F1 indicates that loading C1 is the action most highly afforded for this craft, and a look down the column for all of the other craft, including the Scout, indicates that the affordance for loading this cargo is no higher for any craft other than F1.

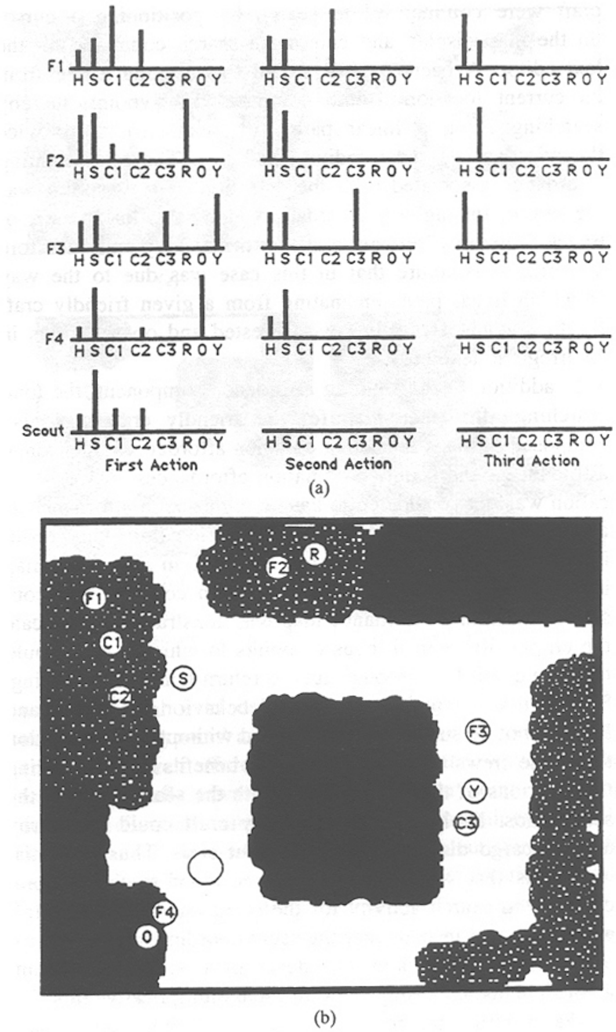

Figure 3. Two Representations of the Same World State: (a) Functional Representation in Terms of Dynamic Affordance Distributions; (b) Representation in Terms of Visual Form

Thus, the model would in this case "decide" to assign the action of loading this piece of cargo to F1.

Given that F1 had been committed in this fashion, the model was then able to determine what the affordances for F1 would be at the time it had completed loading this cargo. This affordance distribution for F1 is shown in the "Second Action" column of distributions. Notice there is no longer any affordance for loading C1 (as this action will have been completed), and now the action of loading the cargo labeled "C2" is most highly afforded. In this case, a plan to load this cargo allowed the model to generate a "Third Action" affordance distribution for F1, in this case indicating that the action of visiting home base "H" would be most highly afforded at that time, due to the opportunity to then score points by unloading two pieces of cargo.

What is absolutely crucial to emphasize, however, is that Figure 3 provides a mere snapshot of what was actually a dynamic system. Just seconds after the situation represented by this snapshot, an event could have occurred that would have resulted in a radical change in the affordance distributions shown (such as the detection of an enemy craft by radar). So, although I have spoken as if the model had committed to plans, these plans actually functioned solely as a resource for prediction, anticipation, and scheduling, rather than as prescriptions for action (cf. Suchman, 1987). The "perceptual" mechanisms in the model, tuned to measure the value of the environmental affordances shown in Figures 2 and 3, could be updated 10 times per second, and the actual process of selecting actions was always determined by the affordances in the "First Action" distribution for all craft. Thus, even though the model would plan when enough environmental and participant-provided constraint on the behavior of the controlled system allowed it to do so, it abandoned many plans as well. A central reason for including a planning horizon in the model was to avoid conflicts among the four craft and the Scout: For example, "knowing" that another craft had a plan to act on some environmental object removed that object from any other craft's agenda, and "knowing" that no other craft's plans did not include acting on some other object increased the affordance for acting on that object for the remaining craft.

The components of the model intended to represent functions performed by internal cognition consisted of the previously mentioned perceptual mechanisms for affordance detection, and also a simple mechanism for combining the affordance measures with priority values keyed to the task payoff structure (e.g., points awarded per type of object processed). Notably, as described in Kirlik et al. (1993) these priority values turned out to be largely unnecessary since an experimental manipulation varying the task payoff structure (emphasizing either loading cargo or engaging enemy craft) by a ratio of 16:1 had *no* measurable effect on the behavior of participants. This finding lent credence to the view that participants' behavior was intimately tailored to the dynamic affordance structure of the *Scout World*, a set of opportunities for action that performers' actions themselves played a role in determining. Due to the fact that behavior involved a continual shaping of the environment, any causal arrow between the two would have to point in both directions (Dewey, 1896, Jagacinski & Flach, 2003). The general disregard of payoff information in favor of exploiting affordances is also consistent with the (or at least my) everyday observation that scattering water bottles around one's home is much more likely to prompt an increase of one's water consumption than any urging by a physician to do so.

Additionally, we manipulated the planning horizon of the model, and found that the variance that resulted was not characteristic of expert-novice differences in human performance. This task apparently demanded less "thinking-ahead" than it did "keeping-in-touch." In support of this view, what *did* turn out to be the most important factor in determining the model's performance, and a plausible explanation for expert-novice differences in this task, was the time required for each perceptual update of the world's affordance structure. As this time grew (from 0.5 s to 2 s), the model (and participants, our validation suggested) got further and further behind in their ability to opportunistically exploit the dynamic set of action opportunities provided by the environment, in a cascading, positive-feedback, fashion. This result highlights that many, if not most, dynamic environments, or at least those we have studied, favor fast but fallible, rather than accurate but slow, methods for profitably conducting one's transactions with the world.

A final observation concerning our affordance-based modeling concerns the oft-stated finding that experts or skilled performers are notoriously unable to verbalize rules or strategies that presumably "underlie" their behavior. When shown a concrete situation or problem, in contrast, these same experts are typically able to report a solution with little effort. This phenomenon is often interpreted using constructs such as "tacit knowledge" (Polanyi, 1966) or automaticity (e.g., Shiffrin & Dumais, 1981). If one does assume, for the sake of discussion, that much procedural knowledge exists in the form of "if *p* then *q*" conditionals or rules, then our *Scout World* modeling provides a quite different explanation of why experts may often be unable to verbalize knowledge. Rather than placing such "if *p* then *q*" rules in the "head" of our model, we instead created perceptual mechanisms that functioned to "see" the world functionally, as affordances, which we interpret as playing the roles of the *p* terms in the "if *p* then *q*" construction. The *q*, on the other hand is the internal response to assessing the world in functional terms, and as such, the "if *p* then *q*" construct is distributed across the boundary of the human-environment system. Or at least this was the case in our computational model.

As such, even if the capability existed to allow our model to introspect and report on its "knowledge," like human experts it could not have verbalized any "if *p* then *q*" rules either, since it contained only the "then *q*" parts of these rules. But if we instead "showed" the model any particular, concrete *Scout World* situation, it would have been able to readily select an intelligent course of action. Perhaps human experts and skilled performers have difficulty reporting such rules for the same reason: At high levels of skill or expertise, these conditionals, considered as knowledge, become distributed across the person-context system, and are thus not fully internal entities (cf. Greeno, 1987, on situated knowledge). Simon (1992) discussed the need to consider not only production rules triggered by symbol structures in working memory but also productions triggered by conditions in the external world to model situated action. Using both types of knowledge representation, Simon noted that, “Productions can implement either situated action or internally planned action, or a mixture of these” (Simon, 1992, p. 125). Our *Scout World* modeling shows that it is certainly possible to computationally model situated action using conditionals in which the *p* elements of “if *p* then *q*” rules exist in the (modeled) external environment rather than in the head. The important point is that computationally modeling the external environment is necessary to give a modeler choice over whether the condition sides of condition-action rules should be located in the model of the head or in the model of the world. Making choices of this type is the essence of modeling cognition whose functionality is distributed boundaries in the human-environment system.

Using Tools & Action to Shape One’s Own Work Environment

In Kirlik (1998a, 1998b) I presented a field study of short-order cooking showing how more skilled cooks used strategies for placing and moving meats to create novel and functionally reliable information sources unavailable to cooks of lesser skill. We observed a variety of different cooks using three different strategies to ensure that each piece of meat (hamburgers) placed on the grill were cooked to the specified degree of doneness (rare, medium, or well).

The simplest (“brute force”) strategy observed involved the cook randomly placing the meats on the grill and using no consistent policy for moving them. As a result, this cook’s external environment contained relatively little functionally relevant information. The second (“position control”) strategy we observed was one where the cook placed meats to be cooked to specified levels at specified locations on the grill. As such, this strategy created functionally relevant perceptual information useful for knowing how well each piece of meat should be cooked, thus eliminating the demand for the cook to keep this information in internal memory.

Under the most sophisticated ("position + velocity control") strategy observed, the cook used both an initial placement strategy as well as a dynamic strategy for moving the meats over time.

Specifically, the cook placed meats to be cooked well done at the rear and rightmost section of the grill. Meats to be cooked medium were placed toward the center of the grill (back to front) and not as far to the right as the meats to be cooked well done. Meats to be cooked rare were placed at the front and center of the grill. Interspersed with his other duties (cooking fries, garnishing plates, etc.), this cook then intermittently "slid" each piece of meat at a relatively fixed rate toward the left border of the grill, flipping them about halfway in their journey across the grill surface. Using this strategy, everything that the cook needed to know about the task was perceptually available from the grill itself, and thus, the meats signaled their own completion when they arrived at the grill's left boundary.

In order to abstract insights from this particular field study that could potentially be applied in other contexts (such as improving the design of frustratingly impenetrable information technology), we decided to model this behavioral situation formally, "to abstract away many of the surface attributes of work context and then define the deep structure of a setting" (Kirsh, 2001, p. 305). To do so, we initially noted that the function of the more sophisticated strategies could perhaps best be understood, and articulated, as creating constraints or correlations to exist between the value of environmental variables that could be directly observed and thus considered "proximal," and otherwise unobservable, covert, or "distal" variables. As such, we were drawn to consider Brunswik's theory of probabilistic functionalism, which represents the environment in terms of exactly these functional, proximal-distal relations (Brunswik, 1956; Hammond & Stewart, 2001; Kirlik, in press) These ideas are articulated within Brunswik's lens model, shown in Figure 4.

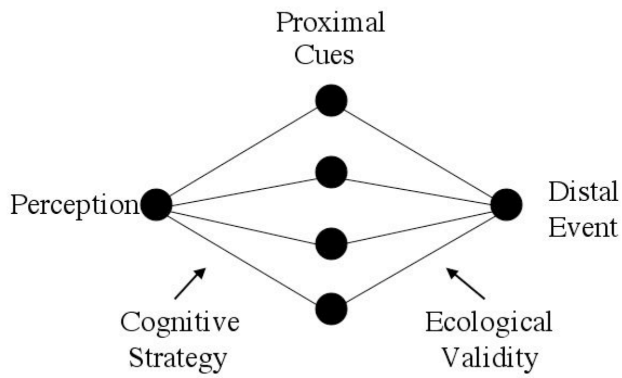

Figure 4. Brunswik's Lens Model of Perception

Brunswik advanced the lens model as a way of portraying perceptual adaptation as a "coming to terms" with the environment, functionally described as probabilistic relations between proximal cues and a distal stimulus. As illustrated in Hammond & Stewart (2001), this model has been quite influential in the study of judgment, where the cues may be the results of medical observations and tests, and the judgment (labeled "Perception" in Figure 4) is the physician's diagnosis about the covert, distal state of a patient (e.g., whether a tumor is malignant or benign). In our judgment research, we have extended this model to dynamic situations (Bisantz et al., 2000), and also to tasks in which cognitive strategies are better described by rules or heuristics rather than by statistical (linear regression based) strategies (Rothrock & Kirlik, 2003). Note that the lens model represents a distributed cognitive system, where half the model represents the external proximal-distal relations to which an agent must adapt to function effectively, and the other half represents the internal strategies or knowledge by which adaptation is achieved.

Considering the cooking case, one deficiency of the lens model should become immediately apparent: In its traditional form it lacks resources for representing the proximal-distal structure of the environment for action, that is, the relation between proximal means and distal ends or goals.

The conceptual precursor to the lens model, originally developed by Tolman & Brunswik (1935), actually did place equal emphasis on proximal-distal functional relations in both the cue-judgment and means-ends realms. As such, we sought to extend the formalization of at least the *environmental* components of the lens model to include both the proximal-distal structure of the world of action, as well as the world of perception and judgment. The structure of the resulting model is shown in Figure 5.

This extended model represents the functional structure of the environment, or what Brunswik termed its "causal texture," in terms of four different classes of variables, as well as any lawful or statistical relationships among them, representing any structure in the manner in which they may covary. The first, [PP,PA] variables are proximal with respect to both perception and action: Given an agent's perceptual and action capacities, their values can be both directly measured and manipulated (in Gibson's terms, they are directly perceptible affordances). [PP,DA] variables can be directly perceived by the agent but cannot be directly manipulated. [DP,PA] variables, on the other hand, can be directly manipulated but cannot be directly perceived. Finally, [DP,DA] variables can be neither directly perceived nor manipulated. Distal inference or manipulation occurs through causal links with proximal variables.

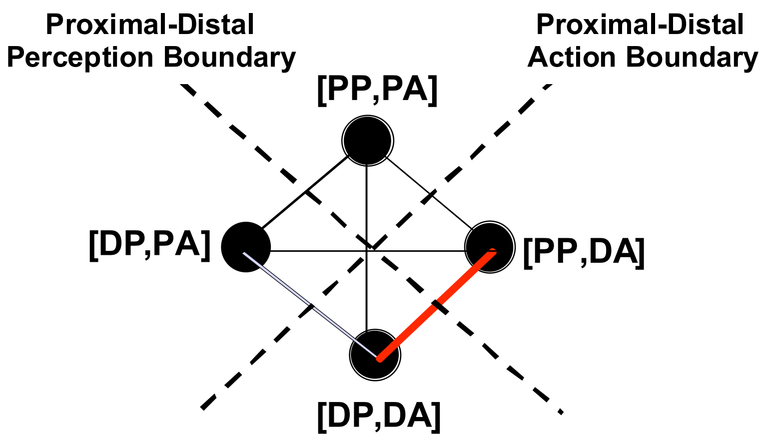

Figure 5. A Functional Model of the Environment for Perception and Action

Note the highlighted link between the [PP,DA] variables and the [DP,DA] variables. These two variable types, and the single link between them, are the only elements of environmental structure that appear in the traditional lens model depicted in Figure 4. All of the additional model components and relations represented in Figure 5 have been added to be able to represent both the functional, perceptual and action structure of the environment in a unified system. See Kirlik (1998a) for a more complete presentation.

To formally analyze the cooking case, we used this model to describe whether each functionally-relevant environmental variable (e.g., the doneness of the underside of a piece of meat) is either proximal (directly perceivable; directly manipulable) or distal (must be inferred; must be manipulated by manipulating intermediary variables), under each of the three cooking strategies observed. Entropy-based measurement (multi-dimensional information theory: see McGill, 1954 for the theory, see Kirlik, 1998a, 2006, for the application to the cooking study), revealed that the most sophisticated cooking strategy rendered the dynamically controlled grill surface not its "own best model" (Brooks, 1991), but rather a fully informative external model of the covert meat cooking process. This perceptible model allowed cooks to offload memory demands to the external world.

Quantitative modeling revealed that the most sophisticated (position + velocity) strategy resulted in by far the greatest amount of variability or entropy in the proximal, perceptual variables in the cook's ecology. This variability, however, was tightly coupled with the values of variables that were covert, or distal to other cooks, and thus this strategy had the function of reducing the uncertainty associated with this cook's distal environment nearly to zero. More generally, we found that knowledge of the demands this workplace task placed on internal cognition would be *underdetermined* without a precise, functional analysis of the proximal and distal status of both perceptual information and affordances, along with a functional analysis of how workers used tools to adaptively shape their own cognitive ecologies.

Modeling the Origins of Taxi Errors at Chicago O'Hare

Figure 6 depicts an out-the-window view of the airport taxi surface in a high-fidelity NASA Ames Research Center simulation of a fogbound Chicago O'Hare airport. The pilot is currently in a position where only one of these yellow lines constitutes the correct route of travel. Taxi navigation errors, and especially errors known as runway incursions, are a serious threat to aviation safety. As such, NASA has pursued both psychological research and technology development in an effort to reduce these errors and mitigate their consequences. In my recent

Figure 6. Simulated View of the Chicago O'Hare Taxi Surface in Foggy Conditions
(Courtesy of NASA Ames Research Center)

collaborative research with Mike Byrne, we completed a computational modeling effort using ACT-R (Anderson, this volume; Anderson & Lebiere, 1998) aimed at understanding why experienced airline flight crews may have committed particular navigation errors in the NASA simulation of taxiing under these foggy conditions (for more detail on the NASA simulation and experiments, see Hooey and Foyle, 2001; Foyle & Hooey, in press; for more detail on the computational modeling, see Byrne and Kirlik, 2005).

Notably, our resulting model was comprised of a dynamic, interactive simulation, not only of pilot cognition, but also of the external, dynamic visual scene, the dynamic taxiway surface, and a model of aircraft (B-767) dynamics. In our task analyses with subject matter experts (working airline captains), we discovered five strategies pilots could have used to make turn-related decisions in the NASA simulation: 1) Accurately remember the set of clearances (directions) provided by air traffic control (ATC) and use signage to follow these directions; 2) Derive the route from a paper map, signage, and what one can remember from the clearance; 3) Turn in the direction of the destination gate; 4) Turn in the direction that reduces the maximum of the X or Y (cockpit-oriented) distance between the aircraft and destination gate; 5) Guess.

We were particularly intrigued by the problem of estimating the functional validity of the two "smart heuristics" (Raab & Gigerenzer, in press; Todd, this volume) involving simply turning in the direction of the destination gate. As such, we provided one of our expert pilots with taxiway charts from all major U.S. airports, and he selected those with which he was most familiar. He then used a highlighter to draw the taxi clearance routes he would likely expect to receive at each of these airports (a total of 258 routes were collected). We then analyzed these routes in terms of their consistency with the two fast and frugal heuristic strategies and found levels of effectiveness as presented in Figure 7.

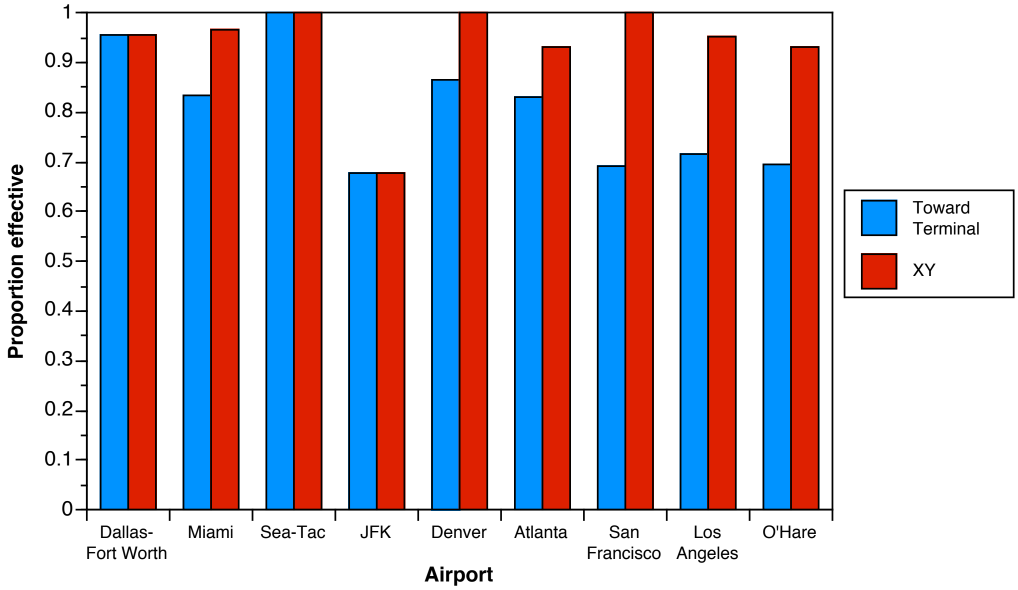

Figure 7. Accuracy of the Two "Fast & Frugal" Heuristics at Nine Major U.S. Airports

We were quite surprised at the effectiveness, or functional validity, of these simple heuristic strategies over such a variety of airports. For example, at "Sea-Tac" (Seattle-Tacoma), these results suggest that a pilot could largely forget the clearance provided by ATC, and simply make a turn toward the destination gate at every decision option, and have ended up fully complying with the clearance that he or she would have most likely been given by ATC. After assembling similar information for all five decision strategies, the ACT-R Monte Carlo analysis of our integrated, functional model resulted in information indicating the frequency with which each of the five strategies would be selected as a function of the decision horizon for each turn in the NASA simulation (see Byrne & Kirlik, 2005),

Specifically, we found that for decision horizons between 2 and 8 seconds, our model predicted that pilots in the NASA experiments would have selected either the "Toward Terminal" or "Minimize XY Distance" heuristics, since within this time interval these heuristics had the highest relative accuracy. Furthermore, an examination of the NASA error data showed that a total of 12 taxi navigation errors were committed. Verbal transcripts indicated that 8 of these errors involved decision making, while the other 4 errors involved flight crews losing track of their location on the airport surface (these "situation awareness" errors were beyond the purview of our model of turn-related decision making).

In support of our functional modeling, every one of the 8 decision errors in the NASA data set involved either an incorrect or premature turn toward the destination gate. Finally, we found that at *every* simulated intersection in which the instructed clearance violated *both* heuristics, at least one decision error was made. In these cases, the otherwise functionally adaptive strategies used by pilots for navigating under low visibility conditions steered them astray, due to atypical structure that defeated their typically rewarded experiential knowledge. Errors did not then result from a general lack of adaptation to the environment, but rather from an overgeneralization of adaptive rules. Generally adaptive decision rules, as measured by their mesh with environmental structure (Todd, this volume) were defeated by ecologically atypical situations.

DISCUSSION

Earlier in this chapter I suggested that modeling interactive behavior and embedded cognition raises theoretical questions that are distinct from the types of theoretical questions that provided the traditional empirical foundation for many cognitive architectures. By "distinct" I meant that some of the theoretical questions that arise when modeling dynamic and interactive tasks are not necessarily reducible in any interesting sense to the questions that motivated the design of these cognitive architectures. I hope that the three modeling examples presented in the previous section are at least somewhat convincing on this point. Each project required us to grapple with problems in cognitive and environmental modeling that I believe to be distinct from the types of questions normally addressed by many cognitive architectures. While one might make the observation that our first two modeling examples, the *Scout World* and short order cooking could have benefited by our use of a cognitive architecture, I would not necessarily disagree. The important point to note is that some sort of detailed functional analyses of those tasks, either those presented or some alternative, would have been required whether or not cognitive architectures were used as the repository for the information gained.

I believe that these examples illustrate the three general points provided in early sections of this chapter on the need to deal head on with theoretical questions arising from the dynamic and interactive nature of embedded cognition. These included the need to model environmental sensitivity to environmental constraints on behavior, the need to model highly context-specific cognitive adaptations, and the need to analyze and model the environment of cognition and behavior in primarily functional terms. While I certainly do not believe that the approach we have taken to these problems represent the final word on these matters, I do hope that these examples have highlighted the need to address them.

REFERENCES

Anderson, J. R., & Lebiere, C. (1998) *The Atomic Components of Thought*. Mahwah, NJ: Lawrence Erlbaum.

Baron, S. (1984). A control theoretic approach to modeling human supervisory control of dynamic systems (pp. 1-48). In W. B. Rouse (Ed.), *Advances in Man-Machine Systems Research, Vol. 1*. Greenwich, CT.

Bisantz, A., Kirlik, A., Gay, P., Phipps, D., Walker, N., and Fisk, A.D., (2000). Modeling and analysis of a dynamic judgment task using a lens model approach. *IEEE Transactions on Systems, Man, and Cybernetics, Vol. 30*, 6, 605-616.

Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence, 47*, 139-159.

Brunswik, E. (1952). The conceptual framework of psychology. In, *International Encyclopedia* *of Unified Science, Vol. 1*, No. 10. (pp. 1-102) Chicago, IL: University of Chicago Press.

Brunswik, E. (1956). *Perception and the Representative Design of Psychological Experiments*. Berkeley, CA: University of California Press.

Byrne, M. & Kirlik, A. (2005). Using computational cognitive modeling to diagnose possible sources of aviation error. *International Journal of Aviation Psychology, 15*(2), 135-155.

Csikszentmihalyi, M. (1993). *Flow: The Psychogy of Optimal Experience*. New York: HarperCollins.

Dewey, J. (1896). The reflex arc in psychology. *Psychological Review, 3*, 357-370.

Donald, M. (1991). *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*. Cambridge, MA: Harvard University Press.

Fajen, B. R. & Turvey, M. T. (2003). Perception, categories, and possibilities for action. *Adaptive Behavior, 11*(4), 279-281.

Foyle, D. & Hooey, R. (in press). *Human Performance Models in Aviation: Surface Operations and Synthetic Vision Systems*. Mahwah, NJ: Erlbaum.

Gibson (1979/1986). *The Ecological Approach to Visual Perception*. Hillsdale, NJ: Erlbaum. (Original work published in 1979).

Gigerenzer, G., Todd, P. M., and the ABC Research Group (1999). *Simple Heuristics that Make us Smart*. New York: Oxford University Press.

Gleick, J. (1992). *Genius: The Life and Science of Richard Feynman*. New York: Pantheon.

Gluck, K. A. & Pew, R. W. (2005) *Modeling Human Behavior with Integrated Cognitive Architectures*. Mahwah, NJ: Erlbaum.

Gray, W. D. & Kirschenbaum, S. S. (2000). Analyzing a novel expertise: An unmarked road. In

J. M. C. Schraagen, S. F. Chipman & V. L. Shalin (Eds.), Cognitive Task Analysis (pp. 275-290). Mahwah, NJ: Lawrence Erlbaum Associates.

Gray, W. D., Schoelles, M. J., & Fu, W. (2000). Modeling a continuous dynamic task. In N. Taatgen & J. Aasman (Eds.), *Proceedings of the Third International Conference on Cognitive Modeling* (pp. 158-168). Veenendal, The Netherlands: Universal Press.

Greeno, J. G. (1987). Situations, mental models, and generative knowledge. In, D. Klahr & K. Kotovsky (Eds.), Complex Information Processing. (pp. 285-316). Hillsdale, NJ: Erlbaum.

Hammond, K. R. & Stewart, T. R. (Eds.). (2001*). The Essential Brunswik*. New York: Oxford University Press.

Hooey, B. L. & Foyle, D. C. (2001). A post-hoc analysis of navigation errors during surface operations. Identification of contributing factors and mitigating strategies. *Proceedings of the 11th Symposium on Aviation Psychology*. Ohio State University, Columbus, OH.

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

Jagacinski, R.J. and Flach, J. (2003). *Control Theory for Humans*. Mahwah, NJ: Erlbaum.

Kirlik, A. (1995) Requirements for psychological models to support design: Toward ecological task analysis. In Flach, J., Hancock, P., Caird, J. and Vicente, K. J. (Eds.), *Global Perspectives on the Ecology of Human-Machine Systems*. (pp. 68-120). Mawhah, NJ: LEA.

Kirlik, A. (1998a). The ecological expert: Acting to create information to guide action. *Fourth Symposium on Human Interaction with Complex Systems*. Dayton, OH: IEEE Computer Society Press: http://computer.org/proceedings/hics/ 8341/83410015abs.htm

Kirlik, A., (1998b). The design of everyday life environments. In W. Bechtel and G. Graham, (Eds.), *A Companion to Cognitive Science*. (pp. 702-712) Oxford: Blackwell.

Kirlik, A. (2006). *Adaptive Perspectives on Human-Technology Interaction: Methods and*

Models for Cognitive Engineering and Human-Computer Interaction. New York: Oxford University Press.

Kirlik, A. & Bisantz, A. M. (1999). Cognition in human-machine systems: Experiential and environmental aspects of adaptation. In P. A. Hancock (Ed.), *Handbook of Perception and Cognition (2nd Ed.): Human Performance and Ergonomics*, pp. 47-68. NY: Academic Press.

Kirlik, A., Miller, R. A., & Jagacinski, R. J. (1993). *IEEE Transactions on Systems, Man, and Cybernetics, Vol. 23*, No. 4, 929-952.

Kirsh, D. (1996). Adapting the environment instead of oneself. *Adaptive Behavior, Vol. 4*, No. 3/4, 415-452.

Kirsh, D. (2001). The context of work. *Human-Computer Interaction, 16*, 305-322.

McGill, W. J. (1954). Multivariate information transmission. *Psychmetrika, 19*(2), 97-116.

Monk, A. (1998). Cyclic interaction: A unitary approach to intention, action and the environment. *Cognition, 68*, 95-110.

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

Newell, A. (1990). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

Newell, A. (1992) Author's response. *Behavioral and Brain Sciences, 15*(3), 464-492.

Payne, J. W. and Bettman, J. (2001). Preferential choice and adaptive strategy use. In G. Gigerenzer & R. Selten (Eds.). (pp. 123-146). *Bounded Rationality: The Adaptive Toolbox.* Cambridge, MA: MIT Press.

Raab, M. & Gigerenzer, G. (in press). Intelligence as smart heuristics. In, R. J. Sternberg & J. E Pretz (Eds.*), Cognition and Intelligence*. New York: Cambridge Univeristy Press.

Rothrock, L. & Kirlik, A. (2003). Inferring rule-based strategies in dynamic judgment tasks: Toward a noncompensatory formulation of the lens model. *IEEE Transactions on Systems,* *Man, and Cybernetics – Part A: Systems and Humans, Vol. 33*, No. 1, 58-72.

Rouse, W. B. (1984). *Advances in Man-Machine Systems Research, Vol. 1*. Greenwich, CT: JAI Press.

Rouse, W. B. (1984). *Advances in Man-Machine Systems Research, Vol. 2*. Greenwich, CT: JAI Press.

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

Shah, K., Rajyaguru, S., St. Amant, R., Ritter, F. E. (2003). Connecting a cognitive model to dynamic gaming environments: Architectural and image processing issues. Proceedings of the Fifth International Conference on Cognitive Modeling (ICCM), pp. 189-194.

Sheridan, T. B. (2002). *Humans and Automation: System Design and Research Issues*. Santa Monica, CA: Human Factors and Ergonomics Society and John Wiley & Sons.

Sheridan, T. B. & Johannsen, G. (1976). *Monitoring Behavior and Supervisory Control*. New York: Plenum Press.

Shiffrin, R. M. & Dumais, S. T. (1981). The development of automatism. In J. R. Anderson, (Ed.), *Cognitive Skills and Their Acquisition*. (pp. 111-140). Hillsdale, NJ: Erlbaum.

Simon, H. A. (1956). Rational choice and the structure of environments. *Psychological Review, 63*, 129-138.

Simon, H. A. (1992). What is an "explanation" of behavior. *Psychological Science, Vol. 3*, No. 3, 150-161.

Suchman, L.A. (1987). *Plans and Situated Actions*. New York. Cambridge University Press.

Tolman, E.C. & Brunswik,E (1935). The organism and the causal texture of the environment. *Psychological Review, 42*, 43-77.

van Geert, P. (2003). Measuring intelligence in a dynamic systems and contextualist framework. In R. J. Sternberg, J. Lautrey, & T. L. Lubart (Eds). *Models of Intelligence: International Perspectives*. (pp. 195 – 212). Washington, DC: American Psychological Association.

Vygotsky, L.S. (1929/1981). The problem of the cultural development of the child, II. *Journal of Genetic Psychology, 36*, 414-434, 1929. The instrumental method in psychology. In J. V. Wertsh (Ed.), *The concept of activity in Soviet psychology*. Armonk, NY: M.E. Sharpe, pp. 134-143, 1981.

Warren, W. H. (1984). Perceiving affordances. Visual guidance of stair climbing. Journal of Experimental Psychology: Human Perception and Performance, 10, 683-703.