

Chapter 8

The Functional Task Environment

*Wayne D. Gray, Hansjörg Neth, and
Michael J. Schoelles*

Although human thought may be possible in those floatation tanks that are used to encourage meditative states, in by far the majority of instances thought occurs in the context of some physical task environment. The physical environment can be as simple as a light and book. It can be as complex as the face of a mountain and the equipment of the climber. It may be as dynamic as the cockpit of an F-16 in supersonic flight and as reactive as a firefight in Iraq or as heated as an argument between lovers.


An emphasis on the environment in cognitive science research is not new. The environment was of prime concern to Simon in his famous “Ant on the Beach” parable (Simon, 1996), in which he warned of the perils of mistaking limits imposed by the environment for limits inherent to human cognition. However, the environment can include an infinity of detail. To be at all useful to understanding human cognition requires a focus on the environment from the perspective of the to-be-accomplished task; that is, it is the task that “allows an environment to be delimited” (Newell & Simon, 1972, p. 55).

The task delimited environment, or more simply the *task environment*, forms the first blade in Newell and Simon’s (1972) oft-quoted scissors analogy:

[AQ1]

Just as a scissors cannot cut paper without two blades, a theory of thinking and problem solving cannot predict behavior unless it encompasses both an analysis of the structure of task environments and an analysis of the limits of rational adaptation to task requirements. (p. 55)

Although the importance of the task environment has been recognized by cognitive science for at least 50 years, it seems fair to say that for most cognitive scientists (especially those working within the experimental psychology tradition), the task environment is something to be rigidly controlled and factored so as to shed light on just one aspect of cognition or one aspect of perception or one aspect of action. Indeed, to dampen further the extent of change during task performance, the other blade of the scissors—human cognition—is also carefully controlled. Many cognitive



studies of both complex (e.g., chess, reading, and so on) as well as very simple tasks (e.g., rapid serial visual presentation, visual search, task switching, and so forth) use expert subjects or train subjects and discard the training trials so that the trial-to-trial operations of the human element remain largely constant.

Unlike critics of contemporary cognitive science research, we do not see these limits on past and current research as an indictment of the failure of cognitive science as a discipline. Rather, we see these limits as a necessary requirement for making advances in our field. Just as it would be unreasonable to ask particle physicists to eschew linear accelerators to pursue their research by studying billiard balls slamming against each other in an actual game of pool, it is also unreasonable to ask cognitive scientists to eschew experimental designs that allow them to isolate and identify the elements and laws of functional cognition.¹ Indeed, these limits on past and some current research have enabled the advances needed for a more integrative approach to cognitive systems.

This chapter introduces the concept of the *functional task environment* (a brief definition is provided in the next section). This concept integrates disparate findings that show important differences between the physical task environment and the ways in which humans perceive, think about, and act on the physical world. The productivity of this concept will be judged by its success at motivating research that leads to the building of integrated models of cognitive systems (Gray, in press). The functional task environment encompasses both blades of Newell and Simon's scissors. Indeed, rather than the metaphor of the scissors, which suggests two structurally independent blades that are used to cut the mental world into small pieces, a metaphor for the functional task environment might be a laser beam that combines "an analysis of the structure of task environments" with "an analysis of the limits of rational adaptation" to provide a strong and focused light onto the operation of a cognitive system that is integrated with the world as well as with perception and action.

The functional task environment is closely related to the main theme of this book: applied attention. Attention operates within the functional task environment and, at the same time shapes it. As we discuss in the next section, the functional task environment is defined over three time spans: evolutionary, life span, and individual tasks. Attention operates both within and on the constraints imposed by each of these time

spans. Indeed, the locus of attention, the cost of shifting attention, and the cost of maintaining attention need to be understood within the context of the functional task environment.

[AQ3]

For this chapter to be successful, by the end its reader will understand the functional task environment to be a concept that draws on and unifies much contemporary cognitive theory. Our further goal is to provide the reader with a new appreciation of the, at times exquisite, adaptation of the functional task environment to the demands of interactive behavior. These adaptations run in both directions: adaptations of the cognitive system to meet the demands of the physical task environment, as well as adaptations of the physical task environment to minimize demands on the cognitive system.

DEFINING THE FUNCTIONAL TASK ENVIRONMENT

The functional task environment emerges from the moment-to-moment intersection of a cognitive agent (human operator) pursuing a particular goal in a particular physical environment. Thus, the functional task environment is mutually constrained by the physical characteristics of the task environment and the functional characteristics of the agent's cognitive system.

It is easy but a mistake to see the contrast as between an objective physical environment and a subjective functional one. The reality is more complex and more interesting. As a first abstraction, the physical task environment is not the same as the complete and objective physical environment in which the task takes place. For example, infrared wavelengths form a part of the objective physical environment. However, because humans are not equipped to perceive these wavelengths, we cannot interact with them and hence they do not form part of our physical task environment.

As a second abstraction, it is clear that many relevant features of the physical task environment are defined for us by the size and shape of our bodies, our strength, and other physiological characteristics. Hence, a pleasant morning swim for a dolphin may be an impossible obstacle for a person. Likewise, a short climb over an overhanging cliff may be an imposing obstacle for a mountaineer.

Finally, as a third abstraction, our physical task environment is defined not simply by reference to the physical environment and our limitations, but by

[AQ2]



other people, human society, cultural artifacts, and inventions. Hence, our buildings, artifacts, infrared detectors, jet skis, and so on change the nature of our physical task environment.

The physical task environment differs considerably from the natural environment but that is not the topic of this chapter. In this chapter we assume a physical task environment and focus our discussion on the factors that control our moment-by-moment interactions within that environment. These factors reflect adaptations to the physical task environment as well as adaptations to the limits of bounded cognition. This division of emphasis between the world and the mind was laid out long ago by Simon (1956, 1992) and has been embraced by thinkers as diverse as Todd (Todd & Gigerenzer, 2000; Todd & Schooler, in press), Anderson (1990, 1991), and Shepard (1990). During the last 15 years, evidence has been amassed that shows just how different our experience of the world—the functional task environment—is from the physical task environment.

As Figure 8.1 is meant to suggest, the functional task environment overlies the physical task environment and the mental system. As such, it is meant as a different level of description than either the physical task environment or the raw properties of the mental system.² The functional task environment encompasses both less and more than the physical task environment. It may not include the objective physical properties of the infrared spectrum, but it does include human inventions such as language and culture that for most of us, most of the time, seem more

real and objective than the infrared spectrum. However, as important as language and culture are to the human experience, we exclude them from the current chapter to focus more clearly on the process of mutual adaptation that governs our immediate interactive behavior.

Within the limits just discussed, we examine the functional adaptations to the physical task environment that have emerged over three very different periods of time:

1. Evolutionary (multiple generations): centuries, millennia
2. Individual life span (learning and development): days, months, years, decades
3. Individual tasks (problem solving, interactive behavior): milliseconds, seconds, minutes, hours

This order of intervals also reflects different parts of Figure 8.1. The next section, “Adaptations over Evolutionary Time,” is illustrated by the bottom of the figure that shows the functional task environment overlaying a subset of the physical task environment. “Adaptations during the Individual Life Span” (a later section) illustrates the top of the figure, which suggests that different tasks and different physical task environments recruit different resources from the mental system. Indeed, this section goes further than this in suggesting that different physical task environments call forth mental resources such as enhanced memory and retrieval resources that are created just for those environments. “Adaptations that Support




FIGURE 8.1. A notional diagram of the functional task environment as a different level of description that overlays a subset of the physical task environment and the mental system.

Interactive Behavior” shows how interactive behavior stitches together the possible elements in the functional task environment into a temporary network of cognition, perception, and action required to do the current task.

the statistical structure of the physical task environment. For purposes of this chapter, the conclusion we draw is that at a very low level of analysis, the functional task environment has already diverged from the physical task environment.

ADAPTATIONS OVER EVOLUTIONARY TIME

The fascination of experimental psychology with perceptual illusions and cognitive illusions (e.g., the “heuristics and biases” research program of Tversky and Kahneman [1974]) has been denounced by critics throughout the decades (Brunswik, 2001/1957; Gigerenzer, 1996) for focusing on the few cases in which humans make faulty judgments rather than the many cases in which human judgments are correct. However, properly regarded, these illusions can be seen as clues to the limits on our experience of the physical world. In this section we discuss three areas—perception, memory, and attention—in which attending to limits has provided a profoundly interesting perspective on the differences between the physical and the functional task environments.

Implications of Perceptual Illusions for the Functional Task Environment

[AQ4]

Purves and associates (2002) point out: “What we see—whether considered in terms of the brightness of objects, their colors or their arrangement in space—is often at odds with the underlying reality measured with photometers, spectrophotometers or rulers” (p. XXX). The fundamental phenomenon is that retinal information is inherently ambiguous. As an example, a small object nearby or a large object far away may generate a line of the same length on the retinal image. Similarly, the same object under the exact same illumination will appear to be a different color depending on its surrounding colors.

Given that our perception of the physical task environment is inherently noisy and ambiguous, what is a mental system to do? The resounding answer from Purves and associates is that the system plays the odds. In a series of fascinating studies (Howe & Purves, 2005a, b; Purves & Lotto, 2003; Purves, Lotto, Williams, Nundy, & Yang, 2001; Purves, Williams, Nundy, & Lotto, 2004; Yang & Purves, 2003), Purves and colleagues show how the wrong answers given to visual illusions tend to reflect the normative answer given

Simple Heuristics that Reflect the Environment in Memory

From a naive perspective it might seem as if a memory system should store and be able to retrieve any experience that we might later want to remember. However, our modern experience with the increasingly large storage requirements of electronic media suggests that even if it were possible to store everything forever, retrieving the right memory when needed would constitute a major problem for any retrieval system. In their article, “Reflections of the Environment in Memory,” Anderson and Schooler (1991) pose what they considered a very basic, but ignored question: “How does a system behave optimally when it is faced with a huge database of items and cannot make all of them instantaneously available?” (p. 396). Their answer—“It would be behaving optimally if it made most available those items that were most likely to be needed”—begot an innovative attempt to define “most likely to be needed” and to determine whether human memory served this function. In this section we briefly review their evidence and arguments and then introduce a recent study by Schooler and others (Schooler & Hertwig, 2005; Todd & Schooler, in press) that shows how this adaptation of memory to its environment can account for one of Gigerenzer and Todd’s “simple heuristics that make us smart” (Gigerenzer & Todd, 1999; Todd & Gigerenzer, 2000).

Characterizing the Demands That the Environment Makes on Memory

Anderson and Schooler (1991) began by trying to define the demands made on memory by the natural environment. They used three sets of data. First were the words that appeared in the *New York Times* headlines for a period of 730 days. They reasoned that headlines posed a demand on the “potential reader of the article to retrieve information about the referent of that word to decide whether this is an article that the reader might want to read” (p. XXX). Second were words from a data set of children’s verbal interactions with adults. “Every time someone

[AQ5]

[AQ6]

[AQ7] says a word to a child, this is a demand on the child to retrieve the word's meaning" (p. XXX). Third was an analysis of senders of e-mail messages to John Anderson over 3.5 years. "The assumption here is that every time Anderson receives a message from a certain person, that is another demand to retrieve some information from Anderson's memory about the sender" (p. XXX).

[AQ8] For each data set they looked at the frequency of occurrence of items for periods of 100 days and asked themselves several questions about day 101. First was the *practice question*: Would the future probability of a word on day 101 reflect its past use in the data set? Second was the *retention question*: Would more remote items occur less likely on day 101? Both questions were confirmed by their analysis. The third question concerned the *spacing effect* for words that appeared twice within a 100-day period: Would their future probability of appearing on day 101 reflect the number of days (space) between the two appearance and the number of days (space) between the second occurrence and day 101? Again, the answer was yes. The data mirrored the human data for studies of massed versus distributed practice. For short lags between the second occurrence and day 101, shorter lags between the first and second occurrences increased the likelihood that the word would occur on day 101. However, for longer lags, use on day 101 increased with the length of the lags between the two occurrences.

[AQ9] Anderson and Schooler (1991) then showed that mathematical functions fitted to the environmental data predicted human data on the effect of practice, forgetting as a function of retention interval, and the spacing effect. They concluded: "This is not a particularly obscure model of the environmental properties of memories. Nonetheless, it turns out these simple assumptions have led to memory characteristics that have confounded psychologists since Ebbinghaus" (p. 408).

***Memory as a Tool in the Adaptive Toolbox:
The Case of the Recognition Heuristic***

US students have been shown to do surprisingly well when asked which of two German cities has the larger population (Gigerenzer & Todd, 1999). Because the US students could only recognize half the German cities in the data set, this success must be based on "ignorance-based reasoning" (Todd & Schooler, in

press). Despite its simplicity, this "recognition heuristic" has been shown to be successful across a wide variety of tasks (Todd & Schooler, in press).

Schooler and others (Schooler & Hertwig, 2005; Todd & Schooler, in press) took the original list of German cities (Goldstein & Gigerenzer, 1999) and determined the probability with which each city was mentioned in the *Chicago Tribune* over a 4767-day period. In their simulation, based on the probability of encountering the city name on each of those 4767 days, the declarative memory element that encoded that city strengthened and decayed according to its frequency and recency of use. The success of this relatively simple model at duplicating the empirical results showed that the recognition heuristic follows from a memory system that reflects the expected utility of information in the environment.

Schooler and Hertwig (2005) did not stop here. They asked the further question of whether forgetting enhanced the accuracy of the recognition heuristic. They studied this by varying ACT-Rs decay parameter and determined that the recognition heuristic worked best with intermediate levels of decay. With too little or too much forgetting, performance declined. An intermediate level maintained a distribution of recognition rates that were highly correlated with the criteria of frequency of mention in the *Chicago Tribune*.

For the purposes of this chapter, this picture of memory reflects an interesting broadening of the functional task environment. In the case of the perceptual phenomenon discussed by Purves and associates, although each new physical stimulus is in some sense new, it is also indistinguishable from thousands or millions of physical stimuli that the cognitive agent has encountered in the past. The perception of an individual stimuli is treated as the perception of a member of a category defined by some physical characteristic.

In contrast, in the case of memory, the function of memory is to recall a particular and, in some sense, distinct item. Hence, the functional task environment for memory reflects an adaptation to a pattern of stimulus occurrences, not to the occurrence of a given stimulus. What is important for memory is the pattern defined by the frequency and recency of an individual item. It is this pattern that predicts whether an individual item will be remembered or forgotten. It is this pattern that enables ignorance-based heuristics to work so well.

[AQ10]

[AQ11]

[AQ12]

**Peering through the Knothole:
Our Functional Representation of
Even Static Physical Task Environments
Is Partial and Dynamic**

The naive view of perception is that we store a high-resolution, surround-sound representation of the physical task environment in our heads. Attending to one part of this representation rather than another seems a simple matter of shifting internal attention. Indeed, when something in the external world changes, it should be a simple matter to compare the new high-resolution representation with the old to detect what has changed.

This caricature of the naive view is unmasked by data (Levin, Momen, Drivdahl, & Simons, 2000) that shows that people consistently and drastically overestimate their ability to detect large changes in visual scenes. Indeed, Findlay and Gilchrist (2003) argue that until recently the view that our representation of our physical task environment is high resolution has dominated, at least implicitly, research on visual perception.

Rensink (2002) shows that the study of *change blindness* emerged out of research on change detection and has now come to define that area. In the mid '90s, researchers (see Findlay and Gilchrist [2003] and Rensink [2002] for a more detailed discussion) studying change detection adopted a paradigm in which changes were made to a visual scene during saccades (we are blind during saccades). They were surprised to discover that in a large number of cases people did not notice the change even when the change consisted of replacing one object for another at the saccade's end point. These changes were in no way subtle. A representative example would be swapping out a picture of a cow for a picture of a motorcycle.

As the work progressed, the changes made by researchers became more and more blatant, yet their subjects still did not notice. In a very dramatic study, Simons and Levin (1998) randomly stopped people on campus to ask for directions. As the victim was gesturing to the questioner, two confederates dressed as workmen and carrying a large door, rudely walked between the questioner and victim. With the door temporarily concealing the questioner from the victim, one workman swapped places with the questioner and continued the conversation as if nothing had happened. Across two replications, only about half the victims noticed the change.

Rensink's work in the laboratory (Rensink, 2000, in press) is almost as blatant and may be more dramatic. Subjects are told to look for changes and for each trial are shown two alternating still shots of the same scene with one major difference. For example, in a scene of Canadian military personnel boarding a transport plane, in one of the two pictures the large engine under one wing is edited out (Fig. 8.2). The engine is fairly large and is toward the center of the picture. The flicker paradigm (Rensink, Oregon, & Clark, 1997) is used in which these pictures alternate continuously until subjects respond that there is or is not a change. Between each of the two pictures, a gray screen is presented that serves to mask the transient abrupt-onset effect. Under these circumstances, although subjects (and large audiences at conference presentations of this work) are staring at the screen, trying to find a change, the change is frustratingly difficult to find. To be clear, the frustration is not the result of the change being small and subtle, but precisely because it is so large and blatant that it is hard to imagine failing to notice it.

The failure to notice a change in the visual display “means that particular piece of information is not part of the internal visual representation” (Hayhoe, 2000, p. XXX). The implications of change blindness for understanding the nature of the functional task environment are profound. In terms of representing the external world, the functional task environment contains less information than the physical task environment. Furthermore, the functional representation of even static physical task environments is partial and dynamic as it “varies from moment to moment in concert with the requirements of the ongoing visual tasks” (Hayhoe, 2000, p. XXX).

[AQ13]

[AQ14]

**Conclusions of Evolutionary Adaptations:
Perception, Memory, and Attention**

The three cases covered in this section tell three different but complementary stories about the relationship of the functional task environment to the physical task environment. Rather than beginning with the assumption that an optimal perceptual system should recognize all angles it encounters without fail, Howe and Purves (2005a) began with the obvious fact that two-dimensional stimuli striking the retina can never fully disambiguate a three-dimensional world. They proceeded to ask about the distribution of angles in the environment and pursued the implications of



FIGURE 8.2. Example of the flicker paradigm. Two still pictures are separated by a gray mask. The pictures alternate until the subject finds the change or concludes that there is no change. See www.psych.ubc.ca/~viscoglab/ for online demonstrations. Printed with permission of R. Rensink.

[AQ15]

the assumption that the perceptual system was adapted to that distribution. Similarly, rather than assuming that an ideal memory would never forget and that all memories are of equal importance, Anderson and Schooler (1991) began by asking how the demands on memory changed with an item's frequency and recency of occurrence. For both perception and memory, in face of uncertainty in the physical task environment, the functional task environment plays the odds.

In the context of the three phenomena discussed in this section, change blindness is the odd man out. Unlike the problems posed by visual illusions and fallible memory, the problem posed by change blindness is a new discovery. Both common sense and common theories of visual attention implicitly assumed that much more of the external world was represented than is actually the case. Indeed, the puzzle here has shifted from explaining what were assumed to be the few cases in which change is not detected, to an accounting of the many cases in which we are functionally blind to changes taking place before our eyes. At least in this regard, theories of visual attention are lagging far behind theories of memory and perception.

ADAPTATIONS DURING THE INDIVIDUAL LIFE SPAN

The functional task environment, as Figure 8.1 suggests, includes portions of the mental system (broadly defined to encompass cognition, perception, and action). Different physical task environments require different mixes of mental processes. Hence, in a very straightforward sense, different physical task environments recruit different mental resources, with the result that the functional task environment differs as a function of the task being performed.

The picture becomes more complicated because with experience in a given physical task environment, people improve at performing the same task. This statement may be the most uncontroversial statement in cognitive science. However, the reasons for this acceleration are not as well understood. Fitts and Posner (1967) talked about three stages of skill acquisition that they called *the cognitive, associative, and autonomous stages*. Anderson (2000) writes of these stages as a transition in knowledge from largely declarative (factual) knowledge that is operated on or interpreted by general-purpose procedures, to a transition stage in which special-purpose procedures begin to

encapsulate the knowledge needed for the task, to an autonomous stage where the set of specialized procedures is complete and can execute the task with little reliance on declarative knowledge and general-purpose, interpretive procedures. Hence, by this standard accounting of skill acquisition, the operators available to the human problem solver change from the slow and error-prone process of acquiring and interpreting declarative knowledge, to the use of specialized routines or procedures.

This sketch of skill acquisition has been generally accepted for the last 40 years. Although modern research has changed our understanding of how these mechanisms work, the story that these are normal changes that use normal mechanisms of cognition holds.

Recently, evidence has accumulated that in addition to these normal changes during skill acquisition, more specialized changes may take place. Some of these changes serve to bypass accepted limits on the time needed to store items in long-term memory, allowing skilled performers to store more information faster. Other changes serve to bypass limits in the number of items that can be retrieved from the canonical 7 ± 2 (Miller, 1956) to 10 or 20 times that amount. Furthermore, neurological evidence is accumulating that, with massive amounts of experience, the brain may change so that different regions become specialized to process different types of material. In this section we discuss life span adaptations of the cognitive system that produce skilled performance in particular functional task environments.

Long-Term Working Memory

Ericsson and Kintsch (Ericsson, 2003; Ericsson & Kintsch, 1995) argue that some types of skilled performance require the use of long-term memory as a type of permanent auxiliary to working memory. They referred to this phenomenon as *long-term working memory* (LTWM) and they argue that their proposal is consistent with all major assumptions regarding long-term and short-term memory. Their disagreement with the standard view focuses on “auxiliary assumptions” regarding the speed of storage and retrieval in long-term memory.

Newell and Simon (1972) presented evidence that it takes between 5 to 10 seconds to store a new and retrievable memory trace, and about 1 second to retrieve an item from long-term memory. Such time scales cannot account for the digit-span experts who

can learn and repeat 50 or more digits (Ericsson & Chase, 1982; Ericsson, Chase, & Faloon, 1980) when read to them at the usual digit-span rate of 1 second per item. Similarly, the best of these experts can memorize and repeat back multiple lists in apparent suspension of the laws of proactive and retroactive interference.

Ericsson and Kintsch (1995) are quite restrained in their claims for the generality of LTWM and are careful to point out that the use of long-term memory as working memory only occurs in the expert’s domain of expertise, where it is “closely tailored to the demands of a specific activity and is an integrated, inseparable part of the skill in performing the activity” (p. 239). They walk through the use of LTWM for the mundane expertise of text comprehension and then provide a quick tour of its use by mental abacus experts, mental calculators (those able to multiple large numbers in their heads), expert waiters (who never write down an order, but always get it right), medical diagnosis, and chess experts. They conclude: “The new contribution we hope to explicate is that reliance on acquired memory skills enables individuals to use [long-term memory] as an efficient extension of [short-term working memory] in particular domains and activities after sufficient practice and training” (p. 211). [AQ16] [AQ17]

Functional Neural Specialization

Individual English letters and Hindu–Arabic numerals have much in common. They are small sets of artificial symbols that share many features and denote abstract entities no one is born knowing. Because of these similarities, Polk and colleagues (Polk & Farah, 1995; Polk, Stallcup, Aguirre, Alsop, D’Esposito, Detre, & Farah, 2002) have become fascinated with the differences between these categories of symbols. In studies of visual search, a pop-out effect occurs when a target and its distracters differ on primitive features such as color. For example, a single red L hidden among multiple green Ts is quickly found. Treisman and Gelade (1980) proposed that such primitive features are processed in specialized modules, and this hypothesis is generally consistent with neuroscientific evidence for spatially segregated cortical areas that process such features (Polk & Farah, 1995). However, the fact that the pop-out effect for red Ls hidden among green Ts may have a neuroscientific basis does not explain why a similar pop-out effect occurs when a given letter is hidden among

numbers (but not when the same letter is hidden among other letters) and vice versa.

Polk and Farah (1995) point out that letters tend to occur with other letters, and numbers tend to occur with other numbers. Hence, if the functional architecture of cognition makes a distinction between letters and numbers, then perhaps the statistical probability of letters co-occurring more with letters than with numbers (and vice versa) interacts with correlation-based learning to lead to maps for letter and digit recognition. They test this hypothesis in two ways.

First, Polk and Farah (1995) found a population of adult subjects for whom the co-occurrence of letters and numbers is much greater than for the population at large. This population was foreign mail sorters at the Philadelphia air mail facility who spend 4 hours each day sorting “Canadian zip codes in which letters and digits occur together (for example, M5S 1A4)” (p. XXX). Testing this population on search time for letters among numbers or vice versa yields a greatly reduced pop-out effect (and longer search times) compared with the control groups. From this they conclude that, “environmental statistics can influence the functional architecture of vision, even in adulthood” (p. 649).

Second, Polk and colleagues (2002) used functional magnetic resonance imaging and found an area of the brain that responds more to letters than to digits. Unlike other areas specialized for visual processing, letters versus digits do not exist in the natural world and can have no evolutionary history. They conclude: “The present finding implies that school-age learning can lead to the creation of new functionally defined brain areas” (p. 154).

Conclusions of Adaptations in Support of Skilled Performance

During the acquisition of skilled performance, repeated experience in a stable physical task environment leads to changes in the functional task environment, which serve to enhance skilled performance. This enhancement includes the recruitment of general-purpose mechanisms that convert declarative knowledge to procedural knowledge. However, it goes beyond such general-purpose mechanisms in that it changes the way the cognitive system processes information and, apparently, changes the neurological architecture of the brain itself.

ADAPTATIONS THAT SUPPORT INTERACTIVE BEHAVIOR

People make adjustments in how they do a task while they are doing it. This statement may seem so obvious that it does not require a demonstration. A light rain on the highway suddenly becomes a torrential downpour. We may quickly turn on our headlights, turn up the speed of the windshield wipers, and reduce our speed. As we do so we may become more attentive to traffic around us, the feel of the road, and so on.

Agre and Shrager (1990) present a detailed analysis of the spontaneous changes one person made while using a photocopier for 4 minutes to make three copies of 17 pages from a book. Their subject came to their study with prior experience using a copying machine but not necessarily prior experience with the particular copying machine used in the study and, almost certainly, no prior experience copying the particular pages from the particular book.

Measuring the time it took the subject to copy each even-odd page pair (e.g., pages 2 and 3, 4 and 5, and so on) they found that work sped up from about 53 seconds for the first pair of pages to around 22 seconds for the final pair. Their analysis makes it clear that this acceleration is not simply the quantitative acceleration of a fixed set of actions such as might be expected by the Fitts and Posner (1967) three-stage analysis of skill acquisition discussed earlier. Rather, qualitative change occurs: As the subject became adapted to the machine, she changed how she handled the book and the machine. Different handling left her hands, body, and the machine in different states with respect to each other. These different states engendered further adaptation and so on.

How these adjustments occur, the role of top-down strategies, tradeoffs between cognitive versus perceptual-motor processes, and so on, is less clear. However, studies that have been done have demonstrated such adjustments over a wide range of tasks that involve a wide range of cognitive, perceptual, and action operations.

In this section we focus on by minute, by second, and by hundred millisecond adjustments to our functional task environment. Some of these adjustments reflect changes in the allocation of mental resources and processes that better adapt us to the physical task environment. Other adjustments reflect changes in the physical task environment to adapt it better to our mental resources and processes. To complicate matters,

[AQ18]

[AQ19]

[AQ20]

these changes occur amid all the differences between the functional and physical task environments that we discussed in the first two sections. Our perceptual system is playing the odds. Our memory system is biased either toward very immediate experiences or toward past experiences that recur with some regularity. Our visual attention does not veridically record the visual environment and may be blind to huge changes. This has the effect of turning stable and static physical task environments into dynamic functional ones. A further complication is that the cognitive resources we can recruit change as we use them and change because we use them. What's a mental system to do? Viewed from this perspective, it may seem to be a miracle that we are able to clothe ourselves and get out the door each morning, let alone productively pursue our careers, lives, and loves.

Fortunately, the mental system is very robust and seems to excel at making adaptations to the physical task environment. In this section we first review the concept of the unit task (Card, Moran, & Newell, 1983) and introduce the concept of interactive routines. In the second section we review a series of studies that show that small manipulations of the physical task environment result in stable and predictable differences in performance and outcomes.

The Unit Task Level

Unit tasks (Card et al., 1983) are subtasks of a larger task that take about 3 to 30 seconds to perform. By definition, “the unit task is fundamentally a control construct, not a task construct” (Card et al., 1983, p. 386). As a control construct, unit tasks are not given by the physical task environment, but result from the interaction of the physical task environment with the control problems faced by the mental system.

The prototypical example of a unit task is the structure imposed by a typist on transcription typing. The physical task environment for transcription typing consists of the dictated speech, a word processor, plus a foot pedal that controls how much of a recording is played back. As speech is typically much faster than skilled typing, the basic problem faced by the typist is how much of the recording to listen to before shutting it off. The efficient typist listens while typing, and the longer he or she listens, the greater the lag between what they are hearing and what they are typing. At some point the typist shuts off the recording and continues to type until she or he can remember no more

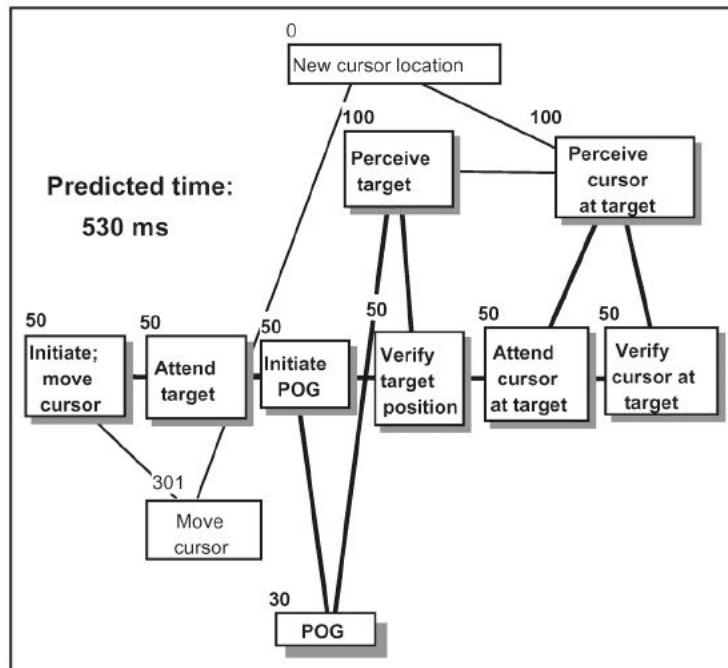
of the recording with certainty. With some experience with the particular speaker and maybe with the particular topic, a skilled transcription typist will minimize the amount of rewind and replay, and maximize the amount typed per unit task. This chopping up of the physical task environment into unit tasks reflects a control process that adjusts performance to the characteristics of the task (the speed of dictation and clarity of speech), to the typist's general typing skill (number of words per minute), as well as to the typist's cognitive, perceptual, and motor limits.

The realm of traditional task analysis (Kirwan & Ainsworth, 1992) lies above the unit task level. As the level of analysis increases from minutes to days, our typist's job may be analyzed into a succession of talks that need transcribing (“transcribe Prof. Wickens' talk, transcribe Prof. Moray's talk”) or even higher level activities that need to be completed (transcribe symposium talks, proofread transcription, send copy of each transcription to the speaker for review). Below the unit task level we would analyze the unit task into a series of interactive routines (Fig. 8.3)—that is, into an activity network of cognitive, perceptual, and action operators (Gray & Boehm–Davis, 2000; Schweickert, Fisher, & Proctor, 2003).

Figure 8.3 provides an example of an interactive routine for moving the cursor to a target location (e.g., menu item, icon, sentence, word, and so forth). The center row shows activities of central cognition that initiate or harvest the activities of other modules. For example, the “initiate move cursor” item initiates a motor command to “move cursor” to a target location. Below the line of central cognitive operators are the motor operators for manual and eye movements. Above the line are visual perceptual operators. Above all these operators is the one box—“new cursor location”—that indicates that a change has been made in the physical task environment.

There are several noteworthy things about interactive routines and their elements for our analysis of the functional task environment. First, interactive routines occur over a time span of 0.333 to 1 second. This level of analysis has been identified by Ballard as the embodiment level (Ballard, Hayhoe, Pook, & Rao, 1997). It is the level at which interactive routines bring together the elements of embodied cognition—namely, cognition, perception, and action.

Second, there is something very fluid about the notion of an interactive routine. The template shown in Figure 8.3 can be instantiated in a number



[AQ21]

FIGURE 8.3. A CPM-GOMS model of the interactive routine required to move a mouse to a predetermined location on a computer screen. Total predicted time is 530 ms. CPM-GOMS is a network modeling technique. In the middle row are cognitive operators with a default execution time of 50 ms each. Above that line are the perceptual operators and below it are the motor operators. The flow of operators is from left to right, with connecting lines indicating dependencies. Within an operator type the dependencies are sequential. However, between operator types the dependencies may be parallel. The numbers above each operator indicate the time, in milliseconds, for that operator to execute. Time is accumulated from left to right along the critical path. The critical path is indicated by bold lines connecting shaded boxes. loc, location; POG, point of gaze. See Gray and Boehm-Davis (2000) for more detailed information.

of different ways. The beginning cursor position, the distance moved, and the size, color, shape of the target all may vary, but the interactive routine would be the same. The operators of an interactive routine seem to encapsulate a low-level control structure that can be used, as in the example provided by Figure 8.3, to move any given cursor, any given distance, to any given object within the physical task environment defined by a computer display.

Adaptations below the Unit Task Level

[AQ22] The development of GOMS (Card et al., 1983) as a tool for cognitive task analysis at the unit task level

coincided with and contributed to the rise of the cognitive engineering movement (Newell & Card, 1985; Norman, 1982, 1986, 1989). This movement strongly implied that because of how changes in design interacted with human cognitive, perceptual, and action resources, small changes in artifacts or task environments could affect both the efficiency and effectiveness of human performance. Given this background, it is perhaps not surprising that some of the first work in this area was done by researchers with strong ties to applied research communities (e.g., Gray & Anderson, 1987; Gray & Orasanu, 1987; John, 1990; John, Rosenbloom, & Newell, 1985; Lohse, 1993; Payne & Green, 1986; Payne, Squibb, & Howes, 1990), who

then turned toward more basic research issues as their work evolved.

In this section we review some of the key studies from the last 15 years that substantiate the claims of the cognitive engineering movement and that have helped to establish basic research in embodied cognition. The reviewed research explores the types of changes that produce effects, as well as the scope of these effects. Speculations about the control mechanisms responsible for these changes are outside the scope of the current chapter. However, such issues have formed the focus of much recent work (Fu & Gray, 2004, in press; Gray & Fu, 2004; Gray, Schoelles, & Sims, 2005; Maloney, Trommershäuser, & Landy, in press; Mozer, Kinoshita, & Shettel, in press).

Do Imposed Changes in the Structuring and Sequencing of Interactive Routines Affect Unit Task Performance?

The utility of analysis below the unit task level for cognitive engineering was shown in the unlikely domain of telephone company toll and assistance operators (TAOs). The study was initiated by a telephone company that was considering replacing its current workstations for TAOs with a proposed new workstation at the cost of approximately \$70 million (in 1990 dollars), but with a projected savings in operating costs of \$12 million per year.

[AQ23]

Using CPM-GOMS (see Fig. 8.3 for an example of this type of analysis), Gray and coworkers (1993) built two models for each of 15 different call types—one for executing the call on the old workstation and one for the proposed workstation. The models tended to predict that the proposed workstation would take expert users longer to operate than the old workstation. Using the phone company's basis for calculating operating costs, this increase in predicted time translated into an increase in annual operating costs of approximately \$2 million per year (rather than the expected savings of \$12 million per year). A 4-month field trial with "live traffic" (i.e., real customers placing real calls) confirmed the predictions of the model both in terms of direction (i.e., proposed workstation slower than old workstation) and magnitude (approximately 1 second per average call slower rather than the expected 4 seconds per call faster).

Although the proposed workstation was in fact faster than the old workstation by all the metrics its designers had touted (e.g., speed to display text on

screen, time to move to and press the most frequently used keys, and so on), the design of the old workstation supported more efficient ordering and interleaving of interactive routines. For example, the old workstation enabled an ordering of operations that tended to put listening to the customer, not operations by the human operator, on the critical path, thereby permitting the TAO to press critical keys and initiate queries to external databases while the customer spoke. Likewise, the old workstation enabled more efficient use of two hands. The proposed workstation required one hand to move between common sequences of two keys, but the old workstation enabled the TAO to move, say, the left hand into position while pressing the first key using the right hand. Although the workstation component of the physical task environment was slower for the old than the proposed workstation, the functional task environment was faster in that it permitted a more efficient coordination of human cognitive, perceptual, and action resources.

Increasing the Cost of Taking an Action and of Accessing Information in the World

Although the study by Gray and coworkers (1993) emphasized the importance of effective interleaving of cognition, perception, and action within a dynamic environment, other studies highlighted the importance of interaction costs to tasks that are often considered more deliberative. For example, Lohse and Johnson (1996) noted changes in the type of decision-making strategies used as the cost of information access changed from a mouse movement and click to an eye movement. Similarly, in what might be the smallest factor manipulated, Ballard and associates (1995) varied the costs of information acquisition from a simple eye movement to a head movement and noted that the shift decreased the number of times that external information was accessed by presumably increasing subjects' reliance on memory.

Other experimenters documented similar trade-offs. For example, across a series of studies, O'Hara and Payne (1998, 1999) varied the cost of making a move in simple tasks such as the "eight puzzle." As costs increased from a simple click on the object to be moved to typing in a string of simple commands, they found that the number of moves made decreased whereas the overall quality of the moves increased. Hence, in the low-cost interface subjects made many

moves before reaching a solution whereas in the high-cost interface a solution was reached in fewer moves.

Imposing a Change in Mental Workload and Cognitive Control

In a series of interesting studies, Carlson and Sohn (Carlson & Sohn, 2000; Sohn & Carlson, 1998) manipulated the control and storage requirements for a series of simple tasks each of which had the same abstract structure. Each task was a four-step task in which the result of each step had to be used as an operand for the next step. For example, if step A yielded 4 as its result and step B was “add (X, 3),” the result of step B would be 7, and 7 would be an operand for step C.

[AQ24] To manipulate cognitive control, Carlson and Sohn (2000) varied whether for each step subjects received the operator first (e.g., “add”) or the missing operand first (e.g., “3”). In a series of studies they consistently found that steps in which the operator was received first were about 200 msec faster than steps in which the operand was received first.

[AQ25] To manipulate mental workload, they varied the storage requirements of their tasks. For example, in Sohn and Carlson’s (1998) spatial navigation task, people might be given the list of all four operators (e.g., left, right, right, up) and required to hold these in memory as they accessed the operands (e.g., 3, 4, 1, 2) one by one to apply each in-the-head operator to an in-the-world operand (i.e., left-3, right-4, right-1, and up-2). The reverse case was also used in which people were required to hold the operands in memory while they accessed the operators one by one. Holding either operands or operators in memory added approximately 600 to 700 msec per step to task performance; however, it did not change the basic finding that people were faster when they accessed operators first than operands first.

[AQ26] Carlson and Sohn’s (2000) results are consistent with those of Gray and colleagues (Gray et al., 1993), which suggested that TAOs could rearrange and interleave interactive routines to shave off several seconds from a 25-second (approximate average) phone call. Both sets of results are consistent with an interpretation of the studies by Lohse and Johnson (1996) and O’Hara and Payne (1998, 1999) that suggest that in the different conditions of their studies, the same set of unit tasks were implemented by different sets of interactive routines. The interactive routines chosen

worked to optimize the cognitive, perceptual, and action resources brought to bear on the physical task environment. When conditions for their application exist, the interactive routine that saves milliseconds is selected and applied—“milliseconds matter” (Gray & Boehm-Davis, 2000)!

Top-down Control of the Sequencing of Cognitive Operations

Because the emphasis in this section is on interactive behavior, most of the examples have an in-the-world component as well as in-the-head ones. However, the concept of interactive routines and of functional task environments extends to the mostly mental world as well. An important question here is whether the selection and operation of mental interactive routines is purely task driven or whether selection can be under conscious or top-down control.

A compelling study of mental interactive routines was provided by Ehrenstein and associates (1997), who required subjects to do two concurrent memory tasks. The memory search task required subjects to hold in memory a set of four, five, or six digits and to indicate after a short delay whether the probe digit was a member of the target set. The arithmetic task presented subjects with a number from four to nine and required them to subtract either a one or two from this number. The search set was presented first, but the probe digit and the arithmetic digit were presented simultaneously. They manipulated top-down cognitive control by instructing subjects to respond to the arithmetic task and then the search task or vice versa.

The analyses were as interesting as they were intricate. Their critical path analysis ruled out alternative interpretations of the data to reveal that subjects were doing one task and then the other, and that the order in which the tasks were performed varied with the instructions that the subjects were given. This study has interesting implications for theories of working memory as well as for theories of cognitive control. For working memory the authors conclude: “Searching a memory set for a displayed item, performing mental arithmetic, and preparing responses to either task all require access to limited working memory processes and appear to be executed sequentially” (Ehrenstein et al., p. 795).

For control of cognition, the authors pit their findings (that the order of processing may be under the top-down or conscious control of the subjects) against

a position that would argue that task-driven or bottom-up processes determine the order of processing. We, however, draw a more general conclusion; namely, that the order of low-level cognitive processes, even those that use the same cognitive resource, is not strictly determined by the physical task environment, but can be influenced by the functional one.

Coordinating and Optimizing the Use of Mental and Motor Operations

Most of the research reviewed in this section has a strong mental or cognitive processing perspective. When motor movements or actions are considered at all, they are viewed primarily as a means of manipulating the cost of a cognitive process. In contrast, a primarily motor movement perspective is provided by Shin and Rosenbaum (2002) who constructed a “nested” aiming arithmetic task that required both perceptual–motor and cognitive processing. They argue that maximum performance in their task required that cognitive processes be coordinated with perceptual–motor ones. Their task required subjects to move a cursor to a small circle (about the size of a typical radio button) on one side of the screen that made task information pop up in an adjacent box. Subjects then were required to move the cursor to the other side of the screen into another small circle to access the next piece of information (which appeared in a box adjacent to that circle). On the first access, subjects saw a single digit. After the first access they saw an operand–operator pair that had to be summed with the current running total. For example, 6, −2, + 4, −5, +1 would equal 4.

Subjects in this study might perform the task by accessing all information before doing any mental calculations. Shin and Rosenbaum (2002) demonstrated that this was not the case. They then argued that to minimize interference, subjects needed to compute a partial sum using the latest operator–operand pair before accessing the next operator–operand pair. However, subjects could meet this constraint in one of two major ways. They could complete all calculations at each step and then move in a strictly serial order (a calculate–aim strategy) or they could let the two processes run on in parallel. They demonstrated that subjects did not do these tasks in serial order, but did the two in parallel.

After concluding that (1) all information is not collected prior to any calculations and (2) the two tasks of calculate and aim are performed in parallel, Shin and

Rosenbaum (2002) then demonstrated that calculation took longer than aiming. This finding suggested that subjects must have adjusted the two processes so that both completed before the cursor entered the next small circle to access the next operator–operand pair. They then ruled out alternative explanations to show that in the nested aiming arithmetic task, movement times slowed down to accommodate the speed of calculations. This is an interesting finding because it once again suggests that low-level processes are being adjusted to meet the demands of a given task environment. Shin and Rosenbaum (2002) instructed their subjects on the basic procedures for doing the task but, unlike Ehrenstein and associates (1997), did not instruct subjects on how to order the subtasks (i.e., how to coordinate aiming and arithmetic). Hence, although it is not clear to what extent consciously adopted top-down strategies contributed to these results, it is clear that these adaptations show an exquisite sensitivity to the demands of the functional task environment.

If Provided the Opportunity to Reduce Memory Load or to Control Scheduling, Can People Take It?

A difficulty in predicting goal-directed behavior is that, although mental processing can flexibly adapt to the physical task environment, when given a chance, humans will alter the physical task environment to reduce the amount of mental processing required. In the studies reviewed in this section, we see an active (if not necessarily deliberate or conscious) adaptation of the physical task environment to enable the use of interactive routines that minimize mental processing. In both directions, the adaptation of the mental to the physical task environment and the adaptation of the physical task environment to resource constraints, the adaptive processes that create the functional task environment act as if milliseconds matter.

An interesting example of this is provided by Kirlik’s (in press) naturalistic investigation of short-order cooks, [who found](#) that experienced cooks utilized the two-dimensional layout of a grill to create a functional task environment that minimizes cognitive workload. The physical position and ordering of a steak on the grill provides the cooks with easy assessment and control of the hard-to-observe variable of doneness.

A similar example is provided by the investigation by Neth and others (Neth & Payne, 2001; Neth,

unpublished doctoral dissertation) of people's spontaneous use of interactive routines and organizing activities for a mental arithmetic task. Neth (unpublished doctoral dissertation) first showed that adults are faster and more reliable when mentally adding numbers involving round results (such as 10, 20, 30) than numbers that do not add up to round sums. For example, if given the three-term problem $4 + 7 + 6$, efficient adders would ignore the serial order and add $4 + 6$ to produce 10 (the round intermediary sum), and then $10 + 7$ to produce the final sum of 17. Those trials during which round intermediate sums were used were generally faster and more accurate than those for which this strategy was not used.

In related work, Neth and Payne (2001) and Neth (unpublished doctoral dissertation) gave subjects long lists of single- and double-digit numbers to sum. Some groups were given access to simple tools like pointers (fingers or cursor) and pens. In general, people who could point at, mark, or move addends (on paper or on a computer screen) adapted their task environment in a variety of ways that increased their performance. Interestingly, the ways in which people used these simple tools varied as a function of their arithmetic expertise. Less expert adders used the paper and pencil to externalize memory for intermediate results and as an aid for performing calculations. In contrast, the better adders made fewer notes but marked off numbers as they were added to facilitate a nonserial adding strategy.

In a similar vein, Cary and Carlson (2001) studied the distribution of working memory demands over internal and external resources. Participants performed a multistep arithmetic task (unlike Neth's task [unpublished doctoral dissertation], the order of these steps was controlled by the experimenter) in which intermediate results were used as operands for subsequent operations. Subjects were encouraged to take notes as needed.

People's distribution and coordination of knowledge in the world versus knowledge in the head varied with experimental conditions. Conditions with a consistent goal structure presumably allowed for more reliable internal strategies and, as predicted by the experimenters, resulted in fewer notes. Likewise, increasing the perceptual–motor costs of note taking also reduced the number of notes—presumably reflecting a willingness to increase mental effort to avoid physical effort. Likewise, note taking decreased when the physical layout of the notes did not corre-

spond to their temporal layout. In this case, keeping track of which note corresponded to which step seems to have required an increase in mental effort (e.g., memory and visual attention) that offset the utility of note taking and increased the likelihood of in-the-head strategies. Finally, Cary and Carlson (2001) found that as people became more expert at the task, they labeled fewer intermediate steps and took fewer notes.

Summary and Conclusions of Interactive Behavior Adaptations

Viewed as a whole, the studies in this section paint a consistent picture: People make adjustments in what they do as they do it. The statement that people tend to distribute working memory resources over internal and external resources seems almost trivially true. This makes it even more surprising how little we currently know about the exact principles and processes by which people spontaneously interact with their physical task environments and spontaneously adapt to situational cost–benefit constraints.

The demands that the functional task environment makes on human cognitive, perceptual, and action operations causes these operations to adapt to each other and to the functional task environment in ways that are defined by various interactive routines. Sometimes these adaptations result in a readjustment that is limited to cognition, perception, and action; other times these adaptations result in changes in the pattern of use of mental versus environmental resources; and sometimes an operator's actions adapt the environment itself, which then may lead to additional adaptations and changes.

Viewed at or below the unit task level, these adaptations do not resemble incremental increases in the speed with which a limited set of processes are executed, but constitute qualitative shifts in the interactive routines used to implement a given unit task. These qualitative shifts sometimes work in direct opposition to simple notions of “speed up with practice” as, for example, in the nested aiming arithmetic task (Shin & Rosenbaum, 2002) in which motor movements slow down to accommodate cognition operations. Likewise, although the swapping often recruits resources from in the world to replace those in the head, these swaps may be temporary. As Cary and Carlson (2001) showed, as experience with a task (especially one with a consistent goal structure) increases, the use of external resources (note taking)

may decrease. In this case it seems as if external resources provide the mental equivalent of “water wings” that support the novice swimmer but, as skill increases, become a hindrance that is removed.

This close look at changes in interactive behavior at or below the unit task level has served to reveal a wide range of adaptations that seem to influence a wide range of higher level goals from serving telephone customers more efficiently (Gray et al., 1993), to the strategies selected for decision making (Lohse & Johnson, 1996), to influencing the nature of planning (O’Hara & Payne, 1998, 1999), to more efficient methods for doing simple arithmetic (Carlson & Sohn, 2000; Cary & Carlson, 2001; Shin & Rosenbaum, 2002; Sohn & Carlson, 1998), to simply following the experimenter’s instructions (Ehrenstein et al., 1997). What has not emerged is any simple and consistent theory for predicting how the physical task environment affects the functional one or vice versa. Unlike the phenomena discussed in earlier sections (with the exception of the change blindness discussion), clear principles have not emerged that can guide us in identifying the key features in the functional task environment to which embodied cognition is adapting. (Although see Gray [in press] for some discussion of the emerging issues on this topic.)

SUMMARY AND CONCLUSIONS

The determinants of human behavior are complex and are often obscured by naive or ill-informed assumptions about the nature of the task environment. It is insufficient to maintain that human behavior is governed by characteristics of the human organism and characteristics of the physical task environment. Instead, we must understand the functional task environment within which the human organism operates.

The twists and turns of the story presented in this chapter may seem very subtle. After offering a broad definition of the functional task environment in the section titled “Defining the Functional Task Environment,” in “Adaptations over Time” we discussed the bottom half of Figure 8.1, which shows the functional task environment overlaying a subset of the physical one. This part of the figure represents adaptations to the physical task environment over evolutionary time. We reviewed recent evidence that suggests that perception was adapted to strike an optimal balance between the realities of using receptors that

distinguish between two dimensions in a three-dimensional world. We also considered evidence that suggested that the memory system functions to provide us with the fastest and most reliable access to those memories that we are most likely to need. We then reviewed evidence that shows that, contrary to our naive experience as well as to older theories of visual perception, we do not have access to a high-resolution representation of all that the eye perceives. Rather, we perceive the world through the knothole of the visual saccade and fixation. Representing an aspect of the visual world is not an automatic consequence of holding our eyes open, but requires moving the knothole to attend to that aspect.

These adaptations result in functional task environments that are so profoundly different from their physical task environments and so seamless that it has literally taken centuries for researchers to notice the discrepancies and provide explanations for how the functional task environment differs from the physical one. Indeed, change blindness is the most recent and arguably most profound discrepancy to be discovered between the physical and functional task environments. Although its pervasiveness has now been well documented, a satisfactory explanation for the factors controlling change blindness has not yet emerged.

As discussed in “Adaptations during the Individual Life Span,” the top half of Figure 8.1 represents the contribution to the functional task environment by the cognitive, perceptual, and action elements of embodied cognition. Much of this part of the functional task environment is defined by normal processes with bounds that experimental psychology has been exploring for more than 100 years. However, modern research is showing that specializations of functional cognition emerge during an individual’s life span throughout the months and years in which skilled practice takes place. Indeed, as the research by Polk and colleagues (2002) shows, these adaptations may result in specialization of the neural architecture.

The nodes and links in Figure 8.1 span the entire functional task environment: the part that overlies the physical task environment as well as the one that overlies the mental system. These nodes and links represent dynamic and temporary adjustments in the structure and performance of individual tasks. Certainly, much of interactive behavior requires the learning of new skills, such as when we learn how to drive a car, touch type, or rappel down a mountain-side. However, although the scope of the changes

discussed in the section titled “Adaptations That Support Interactive Behavior” was wide ranging, none of them seem to involve learning a new skill. Rather, all seemed to involve qualitative changes in how a unit task was implemented or how the physical task environment itself was structured. These qualitative adjustments seem exquisitely sensitive to the cost structure of the functional task environment or to the instructions of the experimenters.

The functional task environment is the playground of immediate behavior. Immediate behavior, when extended in time, is interactive behavior. Interactive behavior presumes a task environment within which behavior interacts. Behavior within a task environment is assumed to be goal directed. In a system with limited resources, it is necessary to use these resources efficiently to accomplish the goal at hand. An efficient use of resources may require restructuring the physical task environment to bring into play more efficient interactive routines.

ACKNOWLEDGMENT

The writing of this chapter was supported by grants from the Air Force Office of Scientific Research (grant no. F49620-03-1-0143) as well as the Office of Naval Research (grant no. N000140310046).

Notes

1. On the subject of making progress by avoiding environmental considerations, we refer the reader to the interesting discussion by Margaret Wilson (2002) on the discovery of the laws governing the properties of hydrogen.
2. We thank Alex Kirlik for pointing out various too-literal interpretations of our earlier descriptions of the functional task environment.

References

- [AQ27] Agre, P. E., & Shrager, J. (1990). Routine evolution as the microgenetic basis of skill acquisition. In *Twelfth annual conference of the Cognitive Science Society* (pp. 694–701). Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, N.J.: Lawrence Erlbaum Associates.
- — — (1991). Is human cognition adaptive? *Behavioral and Brain Sciences*, 14(3), 471–517.
- — — (2000). *Cognitive psychology and its implications* (5th ed.). New York: Worth Publishers.
- — —, & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, 2, 396–408.
- Ballard, D. H., Hayhoe, M. M., & Pelz, J. B. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience*, 7(1), 66–80.
- — —, Hayhoe, M. M., Pook, P. K., & Rao, R. P. N. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, 20(4), 723–742.
- Brunswik, E. (2001/1957). Scope and aspects of the cognitive problem. In K. R. Hammond & T. R. Stewart (Eds.), *The essential Brunswik: Beginnings, explanations, applications* (pp. 300–312). New York: Oxford University Press.
- Card, S. K., Moran, T. P., & Newell, A. (1983). *The psychology of human-computer interaction*. Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Carlson, R. A., & Sohn, M.-H. (2000). Cognitive control of multiple-step routines: Information processing and conscious intentions. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and performance XVIII* (pp. 443–464). Cambridge, Mass.: MIT Press.
- Cary, M., & Carlson, R. A. (2001). Distributing working memory resources during problem solving. *Journal of Experimental Psychology—Learning Memory and Cognition*, 27(3), 836–848.
- Ehrenschein, A., Schweickert, R., Choi, S., & Proctor, R. W. (1997). Scheduling processes in working memory: Instructions control the order of memory search and mental arithmetic. *Quarterly Journal of Experimental Psychology Section A—Human Experimental Psychology*, 50(4), 766–802.
- Ericsson, K. A. (2003). Exceptional memorizers: Made, not born. *Trends in Cognitive Sciences*, 7(6), 233–235.
- — —, & Chase, W. G. (1982). Exceptional memory. *American Scientist*, 70(6), 607–615.
- — —, Chase, W. G., & Faloon, S. (1980). Acquisition of a memory skill. *Science*, 208(4448), 1181–1182.
- — —, & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, 102(2), 211–245.
- Findlay, J. M., & Gilchrist, I. D. (2003). *Active vision: The psychology of looking and series*. New York: Oxford University Press.
- Fitts, P. M., & Posner, M. I. (1967). *Human performance*. Belmont, CA: Brooks Cole.
- Fu, W.-T., & Gray, W. D. (2004). Resolving the paradox of the active user: Stable suboptimal performance in interactive tasks. *Cognitive Science*, 28(6), 901–935.
- — —, & Gray, W. D. (In press). Suboptimal tradeoffs in information seeking. *Cognitive Psychology*. [AQ28]
- Gigerenzer, G. (1996). On narrow norms and vague heuristics: Reply. *Psychological Review*, 103(3), 592–596.
- — —, & Todd, P. M. (Eds.). (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

- [AQ29] Goldstein, D. G., & Gigerenzer, G. (1999). The recognition heuristic: How ignorance makes us smart. In G. Gigerenzer & P. M. Todd (Eds.), *Simple heuristics that make us smart*. New York: Oxford University Press.
- Gray, W. D. (Ed.). (In press). *Integrated models of cognitive systems*. New York: Oxford University Press.
- , & Anderson, J. R. (1987). Change-episodes in coding: When and how do programmers change their code? In: G. M. Olson, S. Sheppard, & E. Soloway (Eds.), *Empirical studies of programmers: Second workshop* (pp. 185–197). Norwood, N.J.: Ablex.
- , & Boehm-Davis, D. A. (2000). Milliseconds matter: An introduction to microstrategies and to their use in describing and predicting interactive behavior. *Journal of Experimental Psychology—Applied*, 6(4), 322–335.
- , & Fu, W.-T. (2004). Soft constraints in interactive behavior: The case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science*, 28(3), 359–382.
- , John, B. E., & Atwood, M. E. (1993). Project Ernestine: Validating a GOMS analysis for predicting and explaining real-world performance. *Human-Computer Interaction*, 8(3), 237–309.
- , & Orasanu, J. (1987). Transfer of learning: Contemporary research and applications. In: S. Cormier & J. Hagman (Eds.), *Transfer of training* (pp. 183–215). San Diego, Calif.: Academic Press.
- , Schoelles, M. J., & Sims, C. R. (2005). Adapting to the task environment: Explorations in expected value. *Cognitive Systems Research*, 6(1), 27–40.
- Hayhoe, M. (2000). Vision using routines: A functional account of vision. *Visual Cognition*, 7(1-3), 43–64.
- Howe, C. Q., & Purves, D. (2005a). Natural-scene geometry predicts the perception of angles and line orientation. *PNAS*, 102(4), 1228–1233.
- , & Purves, D. (2005b). The Muller-Lyer illusion explained by the statistics of image-source relationships. *PNAS*, 102(4), 1234–1239.
- John, B. E. (1990). Extensions of GOMS analyses to expert performance requiring perception of dynamic visual and auditory information. In J. C. Chew & J. Whiteside (Eds.), *ACM CHI'90 conference on human factors in computing systems* (pp. 107–115). New York: ACM Press.
- , Rosenbloom, P. S., & Newell, A. (1985). A theory of stimulus-response compatibility applied to human-computer interaction. In L. Borman & B. Curtis (Eds.), *ACM CHI'85 conference on human factors in computing systems* (pp. 213–219). New York: ACM Press.
- Kirlik, A. (In press). Ecological resources for modeling interactive cognition and behavior. In W. D. Gray (Ed.), *Integrated models of cognitive systems*. New York: Oxford University Press.
- [AQ28] Kirwan, B., & Ainsworth, L. K. (Eds.). (1992). *A guide to task analysis*. Washington, D.C.: Taylor & Francis.
- Levin, D. T., Momen, N., Drivdahl, S. B., & Simons, D. J. (2000). Change blindness blindness: The meta-cognitive error of overestimating change-detection ability. *Visual Cognition*, 7(1-3), 397–412.
- Lohse, G. L. (1993). A cognitive model for understanding graphical perception. *Human-Computer Interaction*, 8(4), 353–388.
- , & Johnson, E. J. (1996). A comparison of two process tracing methods for choice tasks. *Organizational Behavior and Human Decision Processes*, 68(1), 28–43.
- Maloney, L. T., Trommershäuser, J., & Landy, M. S. (In press). Questions without words: A comparison between decision making under risk and movement planning under risk. In W. D. Gray (Ed.), *Integrated models of cognitive systems*. New York: Oxford University Press. [AQ30]
- Miller, G. A. (1956). The magical number 7, plus or minus 2: Some limits on our capacity for processing information. *Psychological Review*, 63(2), 81–97.
- Mozer, M. C., Kinoshita, S., & Shettel, M. (In press). Sequential dependencies in human behavior offer insights into cognitive control. In W. D. Gray (Ed.), *Integrated models of cognitive systems*. New York: Oxford University Press. [AQ31]
- Neth, H., & Payne, S. J. (2001). Addition as interactive problem solving. In J. D. Moore & K. Stenning (Eds.), *Twenty-third annual conference of the Cognitive Science Society* (pp. 698–703). Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Newell, A., & Card, S. K. (1985). The prospects for psychological science in human-computer interaction. *Human-Computer Interaction*, 1(3), 209–242.
- , & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, N.J.: Prentice-Hall.
- (1982). *Steps toward a cognitive engineering: Design rules based on analyses of human error*. Presented at the 1982 conference Human Factors in Computing Systems. Gaithersburg, Md.
- (1986). *Cognitive engineering*. In D. A. Norman & S. W. Draper (Eds.), *User centered system design: New perspectives on human-computer interaction* (pp. 31–61). Hillsdale, N.J.: Lawrence Erlbaum Associates.
- (1989). *The design of everyday things*. New York: DoubleDay.
- , & Payne, S. J. (1998). The effects of operator implementation cost on planfulness of problem solving and learning. *Cognitive Psychology*, 35, 34–70.
- O'Hara, K. P., & Payne, S. J. (1999). Planning and the user interface: The effects of lockout time and error recovery cost. *International Journal of Human-Computer Studies*, 50(1), 41–59.
- Payne, S. J., & Green, T. R. G. (1986). Task-action grammars: A model of the mental representation of task languages. *Human-Computer Interaction*, 2(2), 93–133.
- , Squibb, H. R., & Howes, A. (1990). The nature of device models: The yoked state space hypothesis and some experiments with text editors. *Human-Computer Interaction*, 5(4), 415–444.

- Polk, T. A., & Farah, M. J. (1995). Late experience alters vision. *Nature*, 376(6542), 648–649.
- — —, Stallcup, M., Aguirre, G. K., Alsop, D. C., D'Esposito, M., Detre, J. A., & Farah, M. J. (2002). Neural specialization for letter recognition. *Journal of Cognitive Neuroscience*, 14(2), 145–159.
- Purves, D., & Lotto, R. B. (2003). *Why we see what we do: An empirical theory of vision*. Sunderland, Mass.: Sinauer.
- — —, Lotto, R. B., & Nundy, S. (2002). Why we see what we do: A probabilistic strategy based on past experience explains the remarkable difference between what we see and physical reality. *American Scientist*, 90(3), 236–243.
- — —, Lotto, R. B., Williams, S. M., Nundy, S., & Yang, Z. Y. (2001). Why we see things the way we do: Evidence for a wholly empirical strategy of vision. *Philosophical Transactions of the Royal Society of London Series B—Biological Sciences*, 356(1407), 285–297.
- — —, Williams, S. M., Nundy, S., & Lotto, R. B. (2004). Perceiving the intensity of light. *Psychological Review*, 111(1), 142–158.
- Rensink, R. A. (2000). Visual search for change: A probe into the nature of attentional processing. *Visual Cognition*, 7(1-3), 345–376.
- — — (2002). Change detection. *Annual Review of Psychology*, 53, 245–277.
- — — (In press). The modeling and control of visual perception. In W. D. Gray (Ed.), *Integrated models of cognitive systems*. New York: Oxford University Press.
- — —, Oregan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science*, 8(5), 368–373.
- Schooler, L. J., & Hertwig, R. (2005). How forgetting aids heuristic inference. *Psychological Review*, 112(3), 610–628.
- Schweickert, R., Fisher, D. L., & Proctor, R. W. (2003). Steps toward building mathematical and computer models from cognitive task analyses. *Human Factors*, 45(1), 77–103.
- Shepard, R. N. (1990). *Mind sights*. New York: Freeman.
- Shin, J. C., & Rosenbaum, D. A. (2002). Reaching while calculating: Scheduling of cognitive and perceptual-motor processes. *Journal of Experimental Psychology-General*, 131(2), 206–219.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63, 129–138.
- — — (1992). What is an “explanation” of behavior? *Psychological Science*, 3(3), 150–161.
- — — (1996). *The sciences of the artificial* (3rd ed.). Cambridge, Mass.: MIT Press.
- Simons, D. J., & Levin, D. T. (1998). Failure to detect changes to people during a real-world interaction. *Psychonomic Bulletin & Review*, 5(4), 644–649.
- Sohn, M. H., & Carlson, R. A. (1998). Procedural frameworks for simple arithmetic skills. *Journal of Experimental Psychology—Learning Memory and Cognition*, 24(4), 1052–1067.
- Todd, P. M., & Gigerenzer, G. (2000). Precis of simple heuristics that make us smart. *Behavioral & Brain Sciences*, 23(5), 727–780.
- — —, & Schooler, L. J. (In press). Disintegrated architectures of cognition: The adaptive toolbox for decision making. In: W. D. Gray (Ed.), *Integrated models of cognitive systems*. New York: Oxford University Press.
- Treisman, A. M., & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9(4), 625–636.
- Yang, Z. Y., & Purves, D. (2003). Image/source statistics of surfaces in natural scenes. *Network-Computation in Neural Systems*, 14(3), 371–390.

[AQ32]

[AQ34]