# intellectica

Revue de l'Association pour la Recherche Cognitive


S. Lahlou

D. Kirsh

E. Hutchins

S. Lahlou

A. V. Cicourel

C. Heath, P. Luff, G. Nicholls, D. vom Lehn

W. E. Mackay

M. Zacklad

---


2000/1, n° 30 ISSN n° 0769-4113

## SOMMAIRE

INTRODUCTION

Saadi LAHLOU : *La cognition au travail et ses outils : débordement, révolution, distribution*

*****

David KIRSH : *A Few Thoughts on Cognitive Overload*

Edwin HUTCHINS : *The Cognitive Consequences of Patterns of Information Flow*

Saadi LAHLOU : *Les attracteurs cognitifs et le syndrome du débordement*

Aaron V. CICOUREL : *What counts as data for modeling medical diagnostic reasoning and bureaucratic information processing in the workplace?*

Christian HEATH, Paul LUFF, Gillian NICHOLLS, Dirk vom LEHN : *Textuality and interaction: the collaborative production of new stories*

Wendy E. MACKAY : *Responding to Cognitive Overload: Co-adaptation between Users and Technology*

Manuel ZACKLAD : *La théorie des transactions intellectuelles : une approche gestionnaire et cognitive pour le traitement du COS*

130 F

*Ce numéro a été publié avec le concours du Centre National du Livre, du CNRS et de l'Université de Technologie de Compiègne*

9 770769 411300

*Intellectica*, 2000/1, 30, pp. 53-74

Edwin HUTCHINS*

# The Cognitive Consequences of Patterns of Information Flow

This paper examines the instruments and procedures found on state-of-the-art commercial airline flight decks. It argues that the flight deck as a whole should be viewed as a cognitive system and shows how design decisions may affect patterns of information flow which, in turn, produce cognitive consequences for the flight deck system. Some notable successes and persistent problems are described and a set of design principles that distinguish the successes from failures are presented. Finally, the application of the design principles to long-standing problems illustrates how these principles can lead to new instruments with desirable system properties.

*Keywords*: Distributed cognition, embodied cognition, human computer interaction, aviation instruments, aviation procedures, common-ground understandings, auditory icons.

**Les conséquences cognitives des configurations du flux d'information.** Cet article examine les instruments et procédures que l'on trouve dans les cabines de contrôle des avions commerciaux les plus récents. L'article propose que la cabine de contrôle soit envisagée dans son ensemble comme un système cognitif, et montre comment le design peut influencer les flux d'information qui, à leur tour, produisent des conséquences cognitives pour le système. Quelques succès notables et quelques problèmes persistants sont décrits, et des principes de design qui distinguent les succès et les échecs sont présentés. Enfin, l'application de ces principes à certains problèmes pérennes illustre comment ils peuvent donner naissance à de nouveaux instruments ayant des propriétés souhaitables.

*Mots-clés* : Cognition distribuée, cognition incarnée, interaction homme-machine, instruments de l'aviation, procédures de l'aviation, sens commun, icones auditifs.

* Department of Cognitive Science - University of California, San Diego

## INTRODUCTION

The flight deck of a modern commercial airliner is a complex system consisting of two or more[1] crew and a suite of technological devices. Figure 1 shows the flight deck of the state-of-the-art Boeing 747-400. When everything goes right, all modern flight decks are easy to use. When things go sour, however, automated flight decks provide opportunities for new kinds of problems (Sarter and Woods, 1992). A recent article in Aviation Week cited industry concern over the problem of verifying the safety of complex systems on automated, digital aircraft, stating that the industry must "guard against the kind of incident in which people and the automation seem to mismanage a minor occurrence or non-routine situation into larger trouble." (McKenna, 1998). The design of automated flight deck systems that flight crews find easy to use safely is a challenge in part because this design activity requires a theoretical perspective which can simultaneously cover the interactions of people with each other and with technology.

Figure 1. *Flight deck of the Boeing 747-400*

Aeronautical engineers know how to design reliable systems behind the instrument panel. Behind the mode control panel of

[1] All modern commercial airliner cockpits are designed for two crewmembers. In long-haul operations, however, a third crewmember is present to relieve the primary pilots and to provide information support on takeoff and landing.



Boeing airplanes, for example, there are multiple pathways for information and redundant processors that do automatic error checking by comparing signals (see figure 2). In recent years, this sort of 'system thinking' has continued to develop in computer science. An entire field of computer science called 'distributed artificial intelligence' studies how intelligent behavior arises in multi-processor computer systems. One of the challenges facing aviation today is how to design systems that are reliable on both sides of the instrument panel. Until very recently, design activities and certification procedures focussed on the reliability of hardware and software. While studies of human reliability exist, tools to address the reliability of the system that includes the pilots and technological devices as system components are lacking.

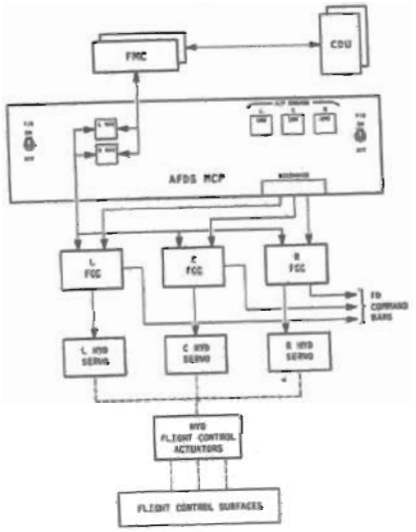


Figure 2. *A system diagram describing what lies behind the mode control panel in a modern airplane.*

In this paper I will introduce some concepts that can be used to understand the flight deck as a system that is composed of two or more pilots and a complex suite of automated devices. As I will try to show, without a theory, we can repeat what seems to work, but we may not know why it worked or how to make it work in novel circumstances. Theory allows us to rise above the particulars of specific situations and makes the roots of success in one setting applicable to other settings.

To the extent that flight decks have been designed with a theory of cognition even implicitly in mind, that theory has usually been the standard information processing theory of cognitive psychology. This view of cognition treats the mind as a logic machine, memory as retrieval from a symbolic data base, and problem solving as a form of logical inference. These central tenets follow from a computational model of the mind and they have implications for the possible roles of the body and social world in explanations of cognitive phenomena. Since cognition is taken to happen in a central processor, "the environment is a problem domain and the body is an input device" (Clark, 1996). The focus on the individual mind as a computational device implies that the individual actor should be the principal unit of analysis for cognition.

## DESIGNING FOR THE TOTAL FLIGHT DECK SYSTEM

Following the psychological theory of human information processing, most attempts to design flight deck systems focus on "the pilot" or on the interaction of "the pilot" with flight deck controls and instrumentation. Even the very best Human-Centered design efforts (for example, Charles Billings' recent book on aviation automation (Billings, 1997) naturally focus on presenting information to "the pilot" rather than to "the flight crew." It is possible to make quite a lot of progress with this sort of theory, but there are many phenomena that are of critical importance to aviation safety that cannot be addressed in a theoretical framework that explains only the properties of individual actors.

In recent years, there have been many calls for design decisions that acknowledge the fact that the crew is a team and that the cognitive system that matters is not the individual pilot or even the pilots. Design must consider the entire flight deck system consisting of two or more people in interaction with each other and now with a complex suite of automated devices. A recent Airline Pilots Association (ALPA) report on "Automated Cockpits and Pilot Expectations" puts the case clearly.

> "The performance of the overall flight crew/flight deck system depends on understanding the total system, its human and automated components and the way these components interact to accomplish the mission. Each design decision should consider overall flight safety and efficiency. Combined flight crew/flight deck system performance is more important than local optimization of the performance of any human or automated component in that system." (ALPA, 1996:5)

During the same period in which industry attention was shifting from the individual pilot to the flight deck system, cognitive theory was changing along similar lines. First, there has been a reconception of the contribution to thinking activity of the body and of structure



in the environment. This is called embodied or ecological cognition (Clark, 1996; Hutchins, 1995a; Lakoff and Johnson, 1999). Second, it has become increasingly clear that cognitive processes may be enacted by two or more persons in interaction with one another. This is the social distribution of cognition (Hutchins, 1995a). Together they comprise a view called "distributed cognition."

These ideas are not unknown in aeronautics, but neither are they fully integrated into human-computer interaction design philosophies. The recent FAA report titled, "The interfaces between flightcrews and modern flight deck systems" (FAA, 1996) has the right wording in its title, referring as it does to "flightcrews" rather than "the pilot". However, a careful reading of the report shows that it is still much easier to think about the cognition of individual pilots than to think about the cognitive properties of the full flight deck system. The question I want to address here is this: How can we actually design for the interactions among the components of the combined flight crew/flight deck system?

There have already been many successful attempts to design for the total cockpit system. Crew Resource Management (CRM), the design of altitude change procedures, and some aspects of lateral navigation displays provide good examples. Still, many aspects of flight deck operations remain problematic. Autoflight mode management is a persistent problem. Why do the successful cases succeed and why do the problems remain unsolved? Until we can answer these questions, we will have to make do with trial and error methods of design.

## SOME SUCCESSES — CRM

Cockpit or Crew Resource Management (CRM) (Wiener, et. al., 1993) provides the very best example of a fully worked out vision of the system view for one aspect of the functioning of the flight deck system. It addresses communication among crewmembers. But there are other aspects that are less well understood and remain largely unrecognized. Hints about a wider perspective on the flight deck systems view come from a variety of sources.

Wiener (1989) conducted observations of crews flying automated airplanes. Among the many phenomena he noted was the fact that captains complained that they had difficulty supervising first officers in the use of the flight management computer system (FMCS) because the control and display unit (CDU) installed at each crew station is difficult to see from the other crew station (figure 3). This is a problem brought on by an interaction between task allocation, display design and the spatial placement of the CDU.

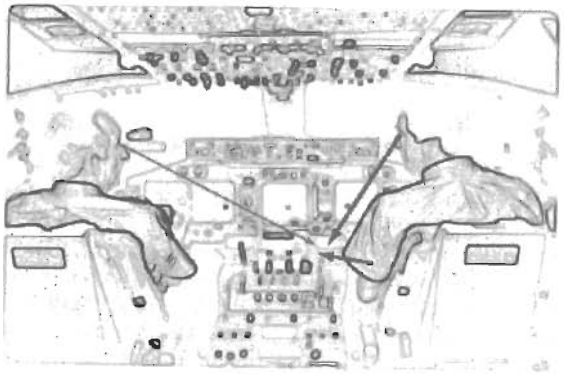

Figure 3. *The flow of information when the First Officer programs the FMS through the CDU. The thin arrow represents the Captain's restricted access to the first officer's actions.*

The difficulty described by the captains concerns the movement of information in the combined flight crew/flight deck system. In particular, the design of the flight deck deprives the captains of information they would like to have. This design interferes with a social distribution of cognition. Good CRM provides a partial solution to this problem by mandating explicit verbal communication of every change made via the CDU, but one is left to wonder how many of the problems solved by CRM were created by design decisions that did not give enough consideration to the patterns of information flow.

## Altitude Change Procedure

In the late 1980s, a major airline instituted a procedure to reduce the number of altitude deviations (Sumwalt, 1991). In this procedure, one pilot sets the target altitude in the mode control panel (MCP) altitude window, announces the altitude, and points to the altitude window until the other pilot reads the altitude shown in the window aloud.

This procedure forces a movement of information so that the altitude is redundantly represented in the minds of the two pilots (figure 4). This procedure improves the social distribution of a cognitive task. Error detection requires a comparison of two or more representations of the same thing. Every pair of arrowheads arriving at a pilot in figure 4 represents an opportunity for error detection. Notice that the modified procedure creates eight opportunities for error detection and correction whereas the standard procedure creates only four. This simple procedural change doubles the number



of opportunities for error detection by constructing an appropriate pattern of information flow.

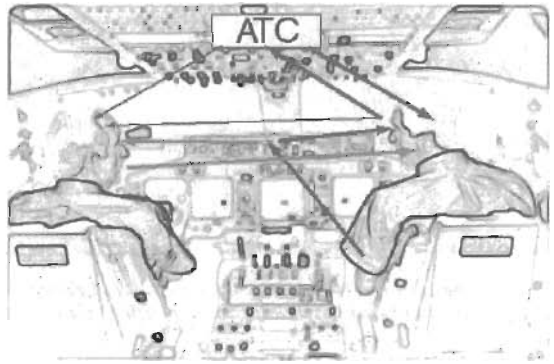


Figure 4. *The flow of information in the altitude change procedure. Heavy arrows indicate mandatory information exchange, light arrows are discretionary, medium arrows represent new mandatory information trajectories introduced by the modified procedure.*

## Descent to an Altitude Restriction

As a final example, consider the task of making a descent to an altitude restriction at a waypoint. In order to perform this maneuver it is necessary to determine a descent rate adequate to satisfy the clearance.

In an older airplane, such as the DC-10, considerable work is required (figure 5). Pilots do this by dividing the distance to the waypoint by present ground speed to get the time to the waypoint. They then divide the difference in altitude between the current altitude and the altitude restriction by the time to the waypoint to determine the vertical rate required.

If the result is a rate that the pilot knows can be safely achieved, he can then adjust pitch, thrust, and drag to produce a vertical rate that will satisfy the clearance. Even allowing for some tricks, such as the fact that Mach number times 10 is approximately the number of miles the airplane covers in a minute, there is a lot of cognitive work to do in order to determine the vertical rate required to meet the restriction. This computation is disembodied decision making. It is done mostly in the head of the pilot and the displays serve only as sources of input to the computation.

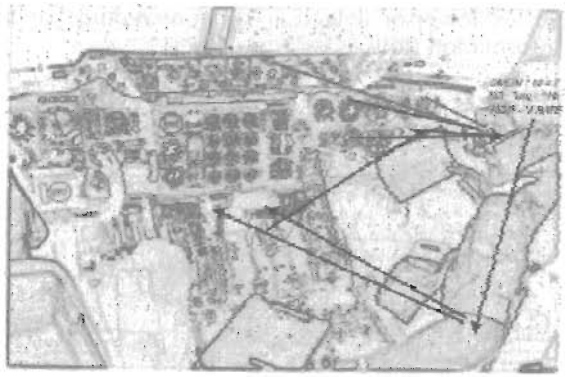

Figure 5. *The flow of information in a non-automated cockpit when making a descent to cross a waypopint at a specified altitude.*

The navigation displays in the current Boeing flight decks offer a feature called the "green arc." This arc appears on the lateral navigation display (moving map) during altitude change maneuvers and shows where, in lateral space, the aircraft will reach the altitude selected in the MCP window based on the present ground speed and vertical speed (figure 6).

To solve the problem in this cockpit, the pilot simply enters the restriction altitude in the altitude window on the MCP and then manipulates the vertical rate until the green arc falls between the airplane and the waypoint. This is a computation, but when a pilot does it using the green arc, it cannot be said that the computation happens entirely in the aircraft computer (because the computer does not contain a representation of the clearance linking the altitude to the waypoint), nor is the computation entirely in the head of a pilot. Rather this computation is performed in the interaction between the pilot, the navigation display, and the way the airplane responds to pilot inputs of thrust, pitch, and drag. Such a performance defies traditional function allocation analyses since this function is not allocated to any single flight deck sub-system. The use of the green arc illustrates the possibility of embodied decision making. The conceptual question of the altitude of the aircraft at the waypoint is represented directly and dynamically in the perceptual properties of the display. With this display it is very easy to monitor progress toward the goal. Further, since the green arc appears on both navigation displays, it facilitates montioring by the pilot not flying and supports the formation of shared understandings about the airplane's situation.



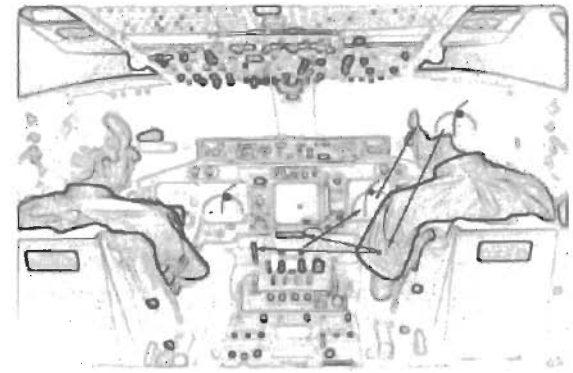

Figure 6. *The flow of information in an automated Boeing cockpit while making a descent to cross a waypoint at a specified altitude.*

Notice that we can apply the same framework to questions concerning pilot interaction with displays that we applied to pilots interacting with one another. In both cases, we ask what information is represented, where and how it is represented, and what patterns of information flow are formed in the activity.

Many design elements that look good in terms of information flow are already present in the cockpit. But the principles of design for information flow have not been clearly articulated. Some of what I have said so far seems to be common sense, but without clear principles it is difficult to fully exploit the potential advantages of well-designed information flow. This is especially true in aviation, where there are complex tradeoffs to be made among elements of common sense that are often in conflict with one another.

## SOME PERSISTENT PROBLEMS

Autoflight mode management is widely recognized as a continuing problem for flight crews. I believe that this is not due to poor engineering of system functionality – the systems are really quite elegant in their operation. Neither is it due entirely to poor training, although the need for improvement here is widely recognized. The problem is located in the user interface to the autoflight functions. Figure 7 illustrates the information flows that are required to make sense of a flight mode annunciation.

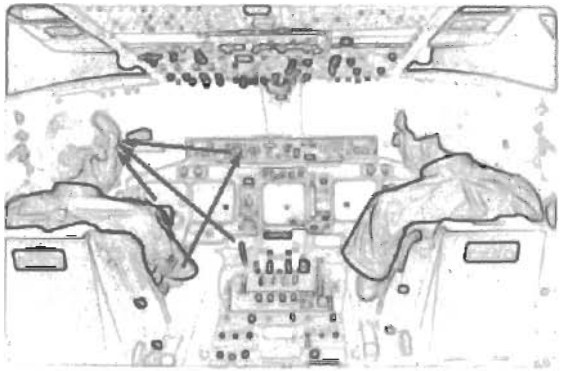

Figure 7. *The flow of information required to interpret the consequences of selecting a high-level autopilot/flight director mode.*

On some models of airplane, when a pilot selects the vertical navigation (VNAV) function on the mode control panel, a flowbar in the VNAV mode select switch on the MCP may be illuminated. This does not mean this mode is engaged, it only means it has been armed for engagement. To know which mode is engaged, one must look to the Flight Mode Annunciator (FMA) which presents strings of text characters that refer to the engaged modes. Under some conditions, VNAV SPD would be displayed.

But what will this mode attempt to do? There are two issues here. First, what targets will VNAV SPD attempt to attain? A speed presumably, and possibly altitudes and other speeds at down-course waypoints as well. To determine what these targets will be, one must consult two different Flight Management Computer System (FMCS) pages via the CDU. Second, how will VNAV SPD attempt to attain those targets? Will it meet the speed target by controlling thrust or pitch? Here, the interface provides very little help. The available indications have been likened to a keyhole through which one peeks at the state of the automation. Determining the functional consequences of mode selection from the character string that names the mode requires complex reasoning. This task could impose a cognitive workload like that encountered in a non-automated airplane while computing the rate of descent required to make a crossing restriction – or perhaps worse if the pilots are unsure how to figure out the functional consequences of the engaged mode. Consider the three descent modes: VNAV PTH, V/S, and MCP SPD. If a pilot is instructed to expedite the descent, what action should he take? Some options are to increase the speed shown in the MCP



speed window, extend the speed brakes, and to click off the autothrottles and increase thrust. Increasing the speed will increase descent rate in VNAV PTH and MCP SPD. Extending the speed brakes will increase descent rate only in MCP SPD. In VNAV PTH it will decrease the descent rate. Adding thrust will increase descent rate in VNAV PTH mode, but will decrease the descent rate in MCP SPD mode. The point is that the functional consequences of alternative modes can be quite different.

The airspeed tapes in many modern cockpits present a less widely recognized problem. To see what the problem is, we need to look first at the instruments that the tapes are intended to replace. Figure 8 shows the airspeed indicator in the Boeing 757/767.

Figure 8. *Airspeed indicator from the Boeing 757/767*

This traditional instrument has some very nice display properties. It maps the abstraction of speed onto the physical space of the speed scale. This mapping permits the pilot to do the conceptual work of reasoning about relationships among speeds with fast robust perceptual processes. With this instrument it is easy to see, even in peripheral vision if one is fast or slow. The full operating speed envelope of the airplane is always visible. The command airspeed bug is just 10 knots wide in the region of the speed scale where approaches are flow. This permits a pilot to know whether the speed is within +/- 5 knots of target with out doing any mental computation at all. The designers of the instrument chose this width for the command airspeed bug so that it would never hide more than one large tic mark at a time[2]. Pilots have discovered that they can put this feature to another use.

[2] Personal communication with the designer of the 757/767 airspeed indicator instrument, September, 1991.

Round-dial airspeed indicators are now being replaced by airspeed tapes. A lot has been said over the past 30 years about the advantages and disadvantages of tape displays. Rather than repeat all that, let me point out that the issues have less to do with the overall format of the display, tape versus round-dial, than they do with the ways the properties of the display can be exploited by pilots while doing their jobs.

Figure 9. *A primary flight display. The airspeed tape is on the left, altitude tape on the right.*

At least one fatal accident has been linked to loss of speed awareness while flying with a digital airspeed tape. Many pilots complain that they are not comfortable using the tapes, but they have difficulty explaining why they are not comfortable. The traditional information processing psychology theory would predict that these displays should be as useful and usable as the electromagnetic displays they replaced. The tapes provide the pilots with an accurate and readable representation of the airplane's airspeed. Common sense tells us that indicating the airspeed must be the function of the airspeed indicator, but a careful examination of what pilots do shows that they get more than an indication of current speed from the airspeed indicator. With such a tape display, it is possible to know if one is fast or slow relative to a target speed, but it is not so easy to know if one is fast or slow relative to the current operating envelope of the airplane. And unlike the electromechanical



displays, it is not possible to know much at all about airspeed from the tape without looking directly at it[3].

In these brief descriptions I have tried to show how attending to the patterns of information flow in the flight deck system can help to explain why some interventions work well and others do not. Let's now look at how the patterns of information flow produce important cognitive properties of flight deck systems. To see how these effects are produced, we need a wider view of information flow, asking what information goes where, when, and under what conditions.

## TRAJECTORIES OF REPRESENTATIONS

### Open interfaces

The feature of design that most influences what information goes where on the flight deck is the physical location of displays, controls, and pilots. Physical locations determine how accessible items are to individual actors and the extent to which interactions with items can be shared by the members of the flight crew. These ideas have been part of flight deck design for decades. Items that are located in shared space and that are easily accessible to both crewmembers, so that each can see what the other may be doing with the item, can be said to have an "open interface." They also support a kind of low-cost monitoring in which each crewmember can maintain awareness of the actions pursued by the other without investing cognitive resources in seeking out information about the actions of the other.

### Consequential Communication

Segal (1995), in a recent paper on design as the choreography of teamwork, provides one of the few attempts to track unintended as well as intended movements of information in a cockpit system. He defines a concept of "consequential communication" as "situations in which engagement in the control task provides operators with information that is relevant to team coordination." Consequential communication makes verbal communication redundant in a situation where redundancy is a desirable property.

### Auditory Icons

Most cockpit displays are visual displays, so spatial placement for vision is the key to their accessibility. Auditory displays are accessible in a different way, so information represented in auditory signals forms different kinds of patterns of information flow. Modern

[3] In fairness, it must be noted that the new airspeed tapes include a number of kinds of useful information that cannot be displayed by electromechanical instruments.

cockpits are equipped with many auditory alerting systems. It seems ironic that over the years so many alarms have been added to cockpits while cockpits have simultaneously been stripped of many useful auditory cues. In the old 727 cockpit, the sound of the trim wheel spinning was an incidental cue that a trim adjustment was in progress. In such a cockpit, a trimming action could become an instance of "consequential communication" because the pilot not flying could know that the pilot flying was adjusting trim without doing anything in particular to seek that information. HCI researchers have recently been exploring a concept they call "auditory icons." Auditory icons are characteristic sounds that accompany actions. It is possible that we should be adding sounds to the cockpit in the form of auditory icons. Describing a groupware interface that uses auditory icons, Gaver and Smith (1990) note, "the sounds perform the dual role of providing feedback to an acting user as well as new information to others. The addition of sound seems to add to the feeling of presence of other users, and to aid in coordinating activities. Sounds that provide confirmatory feedback to one user might well be the only indication to another that certain actions are being performed" (Gaver and Smith, 1990, p. 568). Auditory icons could be a rich source of information about the status of the airplane and especially about what the automation or another pilot is doing to control the airplane. Why not have auditory icons for autoflight mode engagement events?

## REDUNDANT PROCESSING

Patterns of information flow that create multiple representations of the same states of affairs and redundant processing of similar information enhance system robustness. For example, the procedural practice of having both pilots monitor the active air traffic control (ATC) frequency whenever possible creates a pattern of redundant processing and a socially distributed memory for clearances. This gives the flight deck system a property of graceful degradation in the face of local failures. Such patterns of information flow also support error detection, since error detection depends on the comparisons among representations of the same thing developed from different sources or via different processes.

### Common ground

In the context of shared knowledge, patterns of information flow that encourage socially distributed processing of task relevant information lead to the formation of shared expectations about task performance. Shared expectations are the foundation of what is known as common-ground or intersubjective understandings. Intersubjective understanding is a critical component of the



establishment of the meanings of words and actions. It is through intersubjective understanding that even the absence of action can become meaningful in interactions among member of the crew (Hutchins and Klausen, 1998).

For example, when the pilot responsible for communications with ATC fails to respond to a clearance, the absence of a response becomes a very meaningful event for the other crewmember. The meaning given to that absence of action depends on the fact that both crewmembers know what is expected and each knows that the other knows what is expected.

The need for a framework that helps us understand cognitive systems larger than an individual become even more apparent when we consider how cultural factors may affect flight deck performance characteristics. Cultural factors nearly always manifest themselves in the interactions among individuals, and therefore require a theory that reaches beyond the cognitive properties of the individual. Problems of cross-cultural interaction often arise as breakdowns of the "common-ground" due to the lack of sharing of expectations (Perez-Chavez, 1998). Culture is not just something that other people have. It is not something that only appears when people from different backgrounds interact. These circumstances make culture more visible, but culture is present in every moment of every human activity. An American crew flying an American-made airplane in American airspace is a system full of cultural issues. Even if there were only one culture on earth, we would still need a cultural approach. Language is a system of conventional symbolic behaviors. Even if there were only one language on earth, linguistics would still be needed to explain the organization of language behavior. Language gives meaning to words. Similarly, an understanding of culture explains the organization of conventional non-linguistic behavior. Culture gives meaning to action. Cultural issues are essential to the design of safe systems, but they are mostly misunderstood and assumed to be someone else's problem.

CRM and the altitude change procedures described above produce patterns of information flow among pilots that create redundant representations and intersubjective sharing. The green arc on the moving map display illustrates patterns of information flow between pilots and displays and controls that permit conceptual processes to be cleanly mapped into perceptual and motor processes. This permits pilots to use fast, robust perceptual and motor processes to do tasks that would otherwise require brittle reasoning skills.

The problematic cases, autoflight mode management and the airspeed tape, produce patterns of information flow among pilots that fail to create redundant representations and intersubjective sharing.

They create patterns of information flow between pilots and displays and controls that fail to provide clean mappings from conceptual processes to perceptual and motor processes.

These examples highlight the embodied nature of cognition and the social distribution of cognitive processes. They also suggest that features of design may affect patterns of information flow which, in turn, produce cognitive consequences for the flight deck cognitive system.

## AIRSPEED INDICATORS

Let's look forward now. How could this framework help us design displays or procedures for the problematic cases of the airspeed indicators and autoflight mode management?

Consider the round-dial ASI again. The traditional round-dial airspeed indicator provides an example of a display with some useful perceptual properties. In particular, by placing speed bugs on the perimeter of the instruments, pilots are able to imagine conceptually meaningful regions of speed-space. The construction of these conceptually meaningful regions on the face of the instrument permit flightcrews to manage the correspondences between airspeed and required wing configuration by using fast and robust perceptual processes rather than slow conceptual processes (Hutchins, 1995b). As was the case in the use of the green arc, the required computations are in the interaction of the flightcrew with the instrument, rather than inside either the device or the person.

As technology changes, there is always a danger of discarding useful properties that were not recognized in the replaced technology. In their current form, the airspeed tapes that have replaced round-dial instruments in the state-of-the-art cockpits defeat some of the perceptual strategies of pilots. The new instruments offer few perceptually salient cues that pilots can map to their concept of fast/slow in the performance envelope of the airplane. This requires pilots to read the displayed speed as a number and to subject the representation of that speed to further symbolic processing in order to answer the questions that were answered simply by looking at the earlier display.

Boeing's next-generation (NG) cockpits offer a digital round-dial airspeed indicator. This display preserves virtually all of the useful properties of the old electromechanical airspeed indicator and adds in the advantages of digital display (labeled speed bugs for example.)



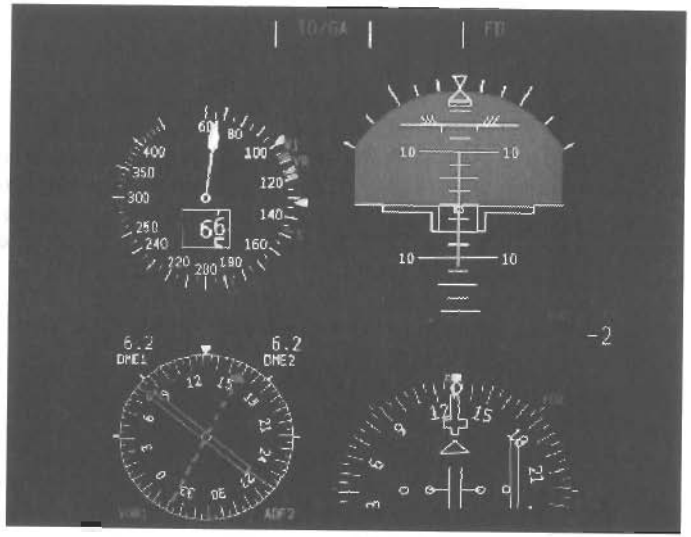


Figure 10. *The primary flight instruments in the Boeing Next-Generation cockpit. The digital recreation of the electromechanical mach/airspeed ndicator (MASI) is at the upper left.*

I predict that pilots will like this instrument, but it doesn't help with the problem of how to redesign the tape. According to the embodied cognition view, two of the most important features of the electromechanical airspeed indicator are first, that its gross physical appearance is different when the airplane is going fast than when it is going slow and second, that the entire operating envelope is always visible. These features permit pilots to develop economical cognitive strategies for interacting with the display. The airspeed tape displays destroy these features. The overall appearance of the display is the same regardless of speed, and usually the limits of the speed envelope are not displayed. The NG airspeed indicator restores these properties by copying the old device. Now that we have an idea what the critical features are, we could add them to the tape display.

Figure 11 illustrates an airspeed tape display that always shows the current operating envelope and permits pilots to know where the current speed is in that envelope. The relative position of the speed box in the display changes with changes in speed. It is perceptually salient and permits the pilot to judge the relative speed of the airplane in peripheral vision.

Figure 11. *Airspeed tape from the Integrated Mode Management Interface.*

The creation of shared spaces that can be partitioned into conceptually meaningful regions supports the development of common-ground understandings, consequential communication, error checking, low cost monitoring, and the substitution of robust perceptual processes for vulnerable conceptual processes.

## THE INTEGRATED MODE MANAGEMENT INTERFACE

We are just beginning to understand what the design criteria could be for a human centered design philosophy that goes beyond the individual human operator and takes seriously the notion that the crew consists of two or more operators interacting with a complex information technology. Let me briefly present an example of an attempt to do such a design.

I have claimed that the perceptual aspects of the display should map cleanly onto the underlying conceptual distinctions made by pilots in doing their work. In order to apply this insight to the problem of autoflight mode management, we would need to know what conceptual features pilots have in mind or should have in mind when they are thinking about autoflight modes. I have been conducting a study to discover the mental models used by pilots to understand autoflight modes – especially vertical navigation modes. To do this I have had pilots make judgements of the similarity of triplets of annunciated mode names. These judgements can be used to reconstruct conceptual distance among modes and plot the results in a concept space as shown in figure 12.

There are two principal dimensions to this space. The modes are implicitly categorized in terms of the source of the performance target (MCP or FMC) and by the type of target (pitch to speed or pitch to path). This indicates that these two distinctions are conceptually salient for pilots.

The integrated mode management interface (IMMI) is a prototype of a graphical user interface to the autoflight system of a state-of-the-art cockpit (Hutchins, 1997). The design of the IMMI is grounded in this study of the conceptual models employed by pilots when they reason about the behavior of the autoflight system. Each of the major



conceptual distinctions used by pilots in reasoning about the autoflight system is represented in the IMMI by a perceptually salient feature. For example, The source of the performance target for an autoflight mode (e.g., MCP or FMCS) is indicated by the physical location of an icon denoting the mode. Icons that appear between the tapes depict modes in which targets are set by the FMCS. Icons that appear outside the tapes refer to modes in which the targets are set by the crew. The type of performance target (e.g., pitch-to-speed mode versus pitch-to-path mode) is indicated by the position of the icon near the speed tape or altitude tape respectively.

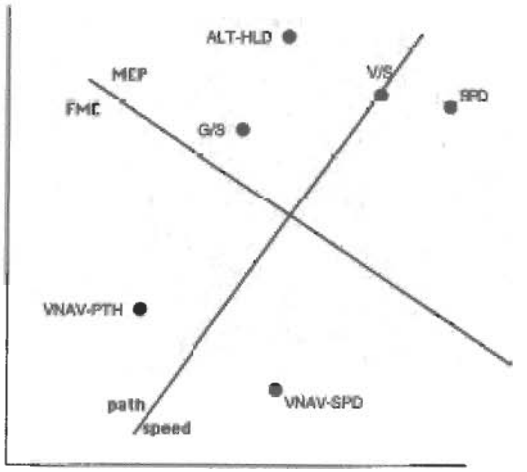


Figure 12. *Conceptual space for the mode names. Distance between each pair of mode names reflects the judged similarity between the modes.*

Engagement status of the mode (e.g., available-for-arming, armed, engaged) is indicated by the color of the icon. The IMMI provides the flight crew with a display in which it is easy to establish conceptually meaningful regions of space. The interface as implemented in the Advanced Concepts Flight Simulator at the NASA-Ames Research Center provides the flightcrew with a shared and accessible control/display system. The representations required to understand the current autoflight mode, to make changes to the mode, and to evaluate the consequences of changes are all located in the same place. This open-interface design promotes patterns of information flow that support the creation of common ground understanding.

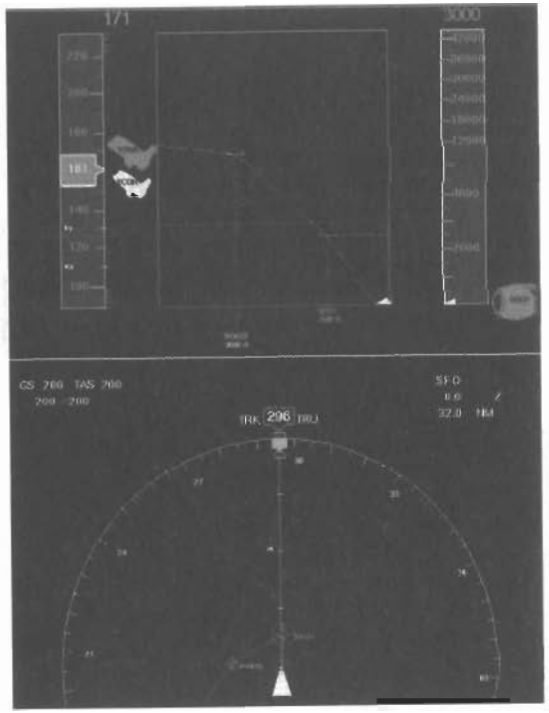

Figure 13. *The integrated Mode Management Interface (IMMI).*

## CONCLUSION

These considerations suggest guidelines for design expressed in terms of the patterns of information flow. In general we should design to create redundant processing across persons, but not within them. This means we should take advantage of opportunities to create consequential communication. We should design for simplicity of processing and shortened trajectories of information within persons. Where possible, we should reduce the number of representations or representational forms that are required to guide the appropriate behavior. Sometimes the cost of changing to a different representation is compensated for by the ease of computation in the new representation. Where possible, the automation should be given the job of transforming representations into forms that are easy for people to use. In general, we should not use automation to make decisions for pilots. Rather we should use automation to present pilots with representations of decision spaces



in which conceptually meaningful distinctions in the decision space are represented by perceptually meaningful distinctions in the appearance or behavior of the display. In such systems, pilots can maintain their engagement with the task of flying and operating the airplane in an embodied rather than disembodied manner. We should be building environments for action in which pilots get to be smart by doing what is easy for them to do. In this way we could bring pilots back into the loop and make them operators in a meaningful world of action rather than supervisory controllers of a fully automated process.

Our current generation of flight deck systems contain many robust and reliable subsystems. I hope the framework I have described here can help to identify patterns in the current operations that work well and suggest ways that future designs can preserve what is good about the current system and perhaps help us to make systems that are even better in the future. Engineers would not think of removing the redundancy of the avionics systems. If we do not understand these patterns on the crew side of the instrument panel, we are at risk of inadvertently disrupting beneficial patterns of information flow.

## *References*


ALPA (1996) Automated Cockpits and Pilot Expectations: A guide for manufacturers and Operators. Unpublished report of the Air Line Pilots Association Automation Team. Herndon Va.

Billings, C. (1997) Aviation Automation. Mahwah, NJ: Lawrence Erlbaum Assoc.

Clark, A. (1996) Being There: Putting brain, body, and world together again. Cambridge, MA: MIT Press.

FAA (1996) The interfaces between flight crews and modern flight deck systems. Federal Aviation Administration human factors team report. June, 1996.

Gaver, W. and Smith, R. (1990) Auditory icons in large-scale collaborative environments. Human-Computer Interaction – INTERACT '90, North-Holland: Elsevier Science Publishers.

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

Hutchins, E. (1995b) How a cockpit remembers its speeds. Cognitive Science, V. 19, pp. 265-288.

Hutchins, E. (1997) The integrated mode management interface. Final report to contract NCC 2-591. NASA-Ames Research Center, Moffett Field, CA. October, 1997.

Hutchins, E. and Klausen, T. (1998) Distributed cognition in an airline cockpit. In Y. Engestrom and D. Middleton, (eds.) Cognition and Communication at Work. Cambridge: Cambridge University Press.

Lakoff, G. and Johnson, M. (1999) Philosophy in the Flesh: The embodied mind and its challenge to Western thought. New York: Basic Books.

McKenna, J. (1998) Industry, FAA struggle to steer agenda. Aviation Week Vol. 148, No.17, p. 61.

Pérez-Chávez, A. (1998) The reconstruction of the social and technological context of the American Airlines 965 accident. Unpublished ms. Wayne State University, Michigan, January 1998.

Sarter, N. and Woods, D. (1992) Pilot interaction with cockpit automation I: Operational experiences with the flight management system. International Journal of Aviation Psychology, V. 2, pp. 303-321.

Segal, L. (1995) Designing team workstations: The choreography of teamwork. In P. Hancock, J. Flach, J. Card, & K. Vicente (Eds.) Local Applications of the Ecological Approach to Human-Machine Systems. Hillsdale, NJ: Lawrence Erlbaum Assoc.

Sumwalt, R. L. (1991) Eliminating Pilot-caused altitude deviations: a human factors approach. Sixth International Symposium on Aviation Psychology, Columbus, Ohio.

Wiener, E. L. (1989) Human factors of advanced technology ("glass-cockpit") transport aircraft. NASA contractor report No. 177528. Moffett Field, CA: NASA-Ames Research Center.

Wiener, E. L., Kanki, B. G., and Helmreich, R. L. (1993) Cockpit Resource Management. New York: Academic Press.

# 9 Cultural Cognition

## The Costs of Failing to See Cognition as a Cultural Process

In this book I have tried to provide a coherent account of cognition and culture as parts of a larger system. This view is not widespread in cognitive science. Yet, there are unnoticed costs in failing to see cognition as part of a cultural process.

### Marginalization of Culture

Early in the development of cognitive science, culture was relegated to a peripheral role. As Gardner (1985) pointed out, culture, history, context, and emotion were all set aside as problems to be addressed after a good understanding of individual cognition had been achieved. It is unfortunate that many anthropologists have encouraged this view by thinking of culture as some collection of things. Tylor (1871) defined culture as "that complex whole which includes knowledge, belief, art, morals, law, custom, and any other capabilities and habits acquired by man as a member of society." Goodenough (1957) gave cognitive anthropology its founding ideational definition of culture: "whatever it is one must know in order to behave appropriately in any of the roles assumed by any member of a society." This view has developed in cognitive anthropology over the years. Attempting to define a role for anthropology in cognitive science, D'Andrade (1981) proposed an intellectual distribution of labor in which psychologists would be responsible for the cognitive processes and anthropologists would be responsible for cognitive content. In this view, culture became simply a pool of ideas that are operated on by cognitive processes. Tylor's definition stresses the acquisition of cultural entities and tries to give a catalog of abilities and artifacts that constitute culture. Goodenough's definition was crucial to the birth of cognitive anthropology, but it and D'Andrade's formulation completely ignore the material aspects of culture. I reject both of these definitions.

Culture is not any collection of things, whether tangible or abstract. Rather, it is a process. It is a human cognitive process that takes place both inside and outside the minds of people. It is the process in which our everyday cultural practices are enacted. I am proposing an integrated view of human cognition in which a major component of culture is a cognitive process (it is also an energy process, but I'm not dealing with that) and cognition is a cultural process.

Anthropologists are also guilty of accepting this marginalization of culture, or even enhancing it, by granting a special place to the powers and limitations of the mind, as if these can be established without reference to culture. Anthropological structuralism tries to read the properties of minds from the structure of public representations. Sahlins (1976) criticizes it as follows: "It would seem ... that the main problem of 'reductionism' besetting modern structuralism has consisted in a mode of discourse which, by giving mind all the powers of 'law' and 'limitation,' has rather placed culture in a position of submission and dependence. The whole vocabulary of 'underlying' laws of the mind accords all force of constraint to the mental side, to which the cultural can only respond, as if the first were the active element and the latter only passive."

Marginalizing culture by reducing it to some collection of ideational contents hides the many ways in which cognition is part of the cultural process. Culture is a process, and the "things" that appear on list-like definitions of culture are residua of the process. Culture is an adaptive process that accumulates partial solutions to frequently encountered problems. It is unfortunate that cognitive science left culture, context, and history to be addressed after the understanding of the individual had matured. The understanding of the individual that has developed without consideration of cultural process is fundamentally flawed. The early researchers in cognitive science placed a bet that the modularity of human cognition would be such that culture, context, and history could be safely ignored at the outset, and then integrated in later. The bet did not pay off. These things are fundamental aspects of human cognition and cannot be comfortably integrated into a perspective that privileges abstract properties of isolated individual minds. Some of what has been done in cognitive science must now be undone so that these things can be brought into the cognitive picture.

### Mistaking the Properties of the System for Those of the Individual

Another cost of failing to see the cultural nature of cognition is that it leads us to make too much of the inside/outside boundary or to assume the primacy of that boundary over other delimitations of cognitive systems.

#### CONSTRUCTION OF PRIMITIVE THOUGHT

A surprising side effect of the heavily drawn inside/outside boundary is that it reinforces the idea that individuals in primitive cultures have primitive minds. The firm drawing of the inside/outside boundary creates the impression that individual minds operate in isolation and encourages us to mistake the properties of complex sociocultural systems for the properties of individual minds. If one believes that technology is the consequence of cognitive capabilities, and if one further believes that the only place to look for the sources of cognitive capabilities is inside individual minds, then observed differences in level of technology between a "technologically advanced" and a "technologically primitive" culture will inevitably be seen as evidence of advanced and primitive minds. Differences in mental capacity seem necessary to account for differences in level of technology. I tried to show in chapters 2–6 that moving the boundaries of the unit of cognitive analysis out beyond the skin reveals other sources of cognitive accomplishment. These other sources are not mysterious, they simply arise from explicable effects that are not entirely internal to the individual.

### Overattribution

Overlooking the cultural nature of cognition has another cost—one that may be the most interesting and far-reaching for the field of cognitive science itself. When one commits to the notion that all intelligence is inside the inside/outside boundary, one is forced to cram inside everything that is required to produce the observed behaviors. Much of cognitive science is an attribution problem. We wish to make assertions about the nature of cognitive processes that we cannot, in general, observe directly. So we make inferences on the basis of indirect evidence instead, and attribute to intelligent systems a set of structures and processes that could have produced the observed evidence. That is a venerable research strategy, and I have no objection to it in principle. However, failing to recognize

the cultural nature of cognitive processes can lead to a misidentification of the boundaries of the system that produced the evidence of intelligence. If we fail to bound the system properly, then we may attribute the right properties to the wrong system or (worse) invent the wrong properties and attribute them to the wrong system. In this attribution game, there has been a tendency to put much more inside than should be there.

## How Cognitive Science Put Symbols in the Head

If there are fundamental deficiencies in the dominant conceptions of cognition in cognitive science, how did that come about?

It is sometimes difficult to say things that are quite simple. The words we must say are simple, but sometimes it takes a lot of work to construct the conceptual framework in which those simple words have the right meanings. There are many possible readings for the sentences I want to write. In the previous chapters I tried to construct some of the conceptual background that will allow me now to say some simple things. However, one hurdle remains. Some of what I have done here departs from the mainstream of cognitive science. And some of the unexamined assumptions of the field make my words unruly. What I want to say cannot be said simply in that framework.

In order to construct a new framework, I will have to deconstruct the old one. In what follows I will give a brief "Official" History of Cognitive Science. This is a history as seen by the proponents of the currently dominant paradigm. I will then reread the history of cognitive science from a sociocultural perspective. In doing this I will identify a number of problems in contemporary cognitive science and attempt to give new meanings to some of the familiar events in its history.

### The "Official" History of Cognitive Science

I begin the official history of cognitive science with a quote from Herbert Simon and Craig Kaplan (1989): "The computer was made in the image of the human."

The ideas on which cognitive science is based are so deeply ingrained in our culture that we can scarcely see how things could be otherwise. The roots of representationalism go back at least to Descartes.

Dreyfus (1992) summarizes the history of Good Old Fashioned Artificial Intelligence (GOFAI) as follows:

*GOFAI is based on the Cartesian idea that all understanding consists in forming and using appropriate symbolic representations. For Descartes, these representations were complex descriptions built up out of primitive ideas or elements. Kant added the important idea that all concepts are rules for relating such elements, and Frege showed that the rules could be formalized so that they could be manipulated without intuition or interpretation.*

The entities that are imagined to be inside the mind are modeled on a particular class of entities that are outside the mind: symbolic representations.

Symbolic logic has a special place in the history of cognitive science. The idea that a computer might be in some way like a person goes back to the formalization of logic and mathematics. In the early years of cognitive science, developments in information theory, neuroscience, psychology, and computer science came to have a synergistic interrelationship. In information theory the notion of a binary digit (bit) as the fundamental unit fit with speculations by McCulloch and Pitts that neurons could be characterized as on/off devices. Thus, the brain might be seen as a digital machine (this turned out to be wrong, but at the time that did not interfere with the developing synergy). Both of these ideas fit well with Turing's work showing that any function that could be explicitly specified could be computed by a class of machine called a universal machine and with his demonstration that the imaginary Turing Machine that operated on a binary code was an example of a universal machine.

The symbol-processing model of cognition has something else going for it as well: "A universal machine can be *programmed* to compute any formally specified function. This extreme plasticity in behavior is one of the reasons why computers have from the very beginning been viewed as artifacts that might be capable of exhibiting intelligence." (Pylyshyn 1989: 54) This was an essential component of the history of the field. Referring to the human cognitive architecture, Newell et al. (1989: 103) say that "the central function of the architecture is to support a system capable of universal computation." By choosing a formalism that is capable of any specifiable computation, the early theorists were surely casting a wide enough net to capture human cognition, whatever it might

turn out to be. It seemed that the only viable challenge to this view would be a demonstration that human cognition might not be formally specifiable. There are many varieties of systems capable of universal computation. Newell and his colleagues and most others in the classical camp have taken what is called a "physical symbol system" as the primary architecture of human cognition. "A physical symbol system is an instance of a universal machine. Thus the symbol system hypothesis implies that intelligence will be realized by a universal computer." (Newell and Simon 1990) Newell and Simon (ibid.) define a physical symbol system this way:

*A physical symbol system consists of a set of entities, called symbols, which are physical patterns that can occur as components of another type of entity called an expression (or symbol structure). Thus a symbol structure is composed of a number of instances (or tokens) of symbols related in some physical way (such as one token being next to another). At any instant of time the system will contain a collection of these symbol structures. Beside these structures, the system also contains a collection of processes that operate on expressions to produce other expressions: processes of creation, modification, reproduction, and destruction. A physical symbol system is a machine that produces through time an evolving collection of symbol structures. Such a system exists in a world of objects wider than just these symbolic expressions themselves.*

According to Pylyshyn (1989), the notion of mechanism that underlies the classical concept of cognition is "concerned only with abstractly defined operations such as storing, retrieving, and altering tokens of symbolic codes."

Simon and Kaplan (1989) cite the Logic Theorist of Newell and Simon (1956) as an example of abstract intelligence and note the role of psychological research in its design:

*The earliest artificial intelligence programs (for example, the Logic Theorist (Newell and Simon 1956)) are perhaps best viewed as models of abstract intelligence; but nonetheless their design was informed by psychological research on memory and problem solving—note, for example, the use of associative structures in list-processing programming languages and subsequently the frequent use of means-ends analysis for inference.*

By embodying the growing knowledge of human information-processing psychology in computer programs, the early researchers

were able to express their theories about cognition as working models that, in many cases, were capable of actually reproducing many important aspects of the behavior of human subjects.

Artificial intelligence (AI) and information-processing psychology thus have a synergistic relationship to each other. Information-processing psychology investigates humans as information processors via the computational metaphor of mind, while AI investigates machine implementations of intelligent processes. The operation of machines that are purportedly built in the image of humans is believed to shed light on natural human intelligence. Since the properties of abstract systems of intelligence are not dependent on the implementational details of the machines on which they run, intelligence in general (in addition to specifically human intelligence) can be investigated with this technology. The hope is that these traditions will continue to synergistically feed each other. In the most optimistic versions of the story, AI and information-processing psychology are the principal motors of scientific progress in cognitive science.

### An Alternative History of Cognitive Science

Let us back up and examine the history of computers a bit more. The digital computer is a physical device that can support a mechanized version of a formal system. And it is this capacity that makes it a potential model of intelligence. Understanding computers requires an understanding of formal systems. We know that formal systems go back several thousand years in the history of our species. I do not know when the formal aspects of formal systems were first understood. I suspect that real understanding of the formal aspects of formal systems did not come until the revolutionary work on mathematics and logic at the beginning of this century that was critical to the foundation of cognitive science. Formal systems themselves are much older than our explicit understanding of them. The first arithmetic systems are at least 3000 years old, so we may take that as a minimum age of formal systems in the human experience. The idea of a formal system is that there is some world of phenomena, and some way to encode the phenomena as symbols. The symbols are manipulated by reference to their form only. We do not interpret the meanings of the symbols while they are being manipulated. The manipulation of the symbols results in some other symbolic expression. Finally, we may interpret a newly

created string of symbols as meaning something about the world of phenomena.

Being able to find sets of syntactic manipulations of symbols that preserve this relationship so that we can reinterpret symbolic expressions into the world is of paramount importance. As Pylyshyn (1989) says: "One might ask how it is possible for symbolic expressions and rules to keep maintaining their semantic interpretation, to keep the semantics of the expressions coherent. It is one of the important discoveries of formal logic that one can specify rules that operate on symbolic expressions in such a way that the sequence of expressions always corresponds to a proof." If we built the right formal system, we could now describe states of affairs in the world that would have been impossible or impractical to observe directly. Such a state of affairs might be something in the future, which we cannot observe directly, but which can be predicted. I consider the mastery of formal systems to be the key to modern civilization. This is a very, very powerful idea.

The system of ship navigation that I have presented in this book is based on formal manipulations of numbers and of the symbols and lines drawn on a chart. It is a system that exploits the powerful idea of formal operations in many ways. But not all the representations that are processed to produce the computational properties of this system are inside the heads of the quartermasters. Many of them are in the culturally constituted material environment that the quartermasters share with and produce for each other.

Now, here is what I think happened. It was discovered that it is possible to build machines that can manipulate symbols. The computer is nothing more than an automated symbol manipulator. And through symbol manipulation one can not only do things we think of as intelligent, like solving logical proofs or playing chess; we know for a fact that through symbol manipulation of a certain type it is possible to compute any function that can be explicitly specified. So, in principle, the computer could be an intelligent system. The mechanical computers conceived by Charles Babbage to solve the problem of unreliability in human compilers of mathematical and navigational tables were seen by his admirers to have replaced the brain: "The wondrous pulp and fibre of the brain had been substituted by brass and iron; [Babbage] had taught wheelwork to think" (H. W. Buxton, cited in Swade 1993). Of course, a century later it would be vacuum tubes that created the "electronic brain."

But something got lost in this move. The origin myths of cognitive science place the seminal insights of Alan Turing in his observations of his own actions. Dennett (1991) describes the context of Turing's discoveries:

*He was thinking self-consciously and introspectively about just how he, a mathematician, went about solving mathematical problems or performing computations, and he took the important step of trying to break down the sequence of his mental acts into their primitive components. "What do I do," he must have asked himself, "when I perform a computation? Well, first I ask myself which rule applies, and then I apply the rule, and then write down the result, and then I look at the result, and then I ask myself what to do next, and. . . ."*

Originally, the model cognitive system was a person actually doing the manipulation of the symbols with his or her hands and eyes. The mathematician or logician was visually and manually interacting with a material world. A person is interacting with the symbols and that interaction does something computational. This is a case of manual manipulation of symbols.

Notice that when the symbols are in the environment of the human, and the human is manipulating the symbols, the cognitive properties of the human are not the same as the properties of the system that is made up of the human in interaction with these symbols. The properties of the human in interaction with the symbols produce some kind of computation. But that does not mean that that computation is happening inside the person's head.

John Searle's "Chinese room" thought experiment provides a good example of this effect. Imagine a room inside of which sits the philosopher Searle. Chinese people come up to the room and push strings of Chinese characters through a slot in the door. Searle slips back other strings of characters, which the Chinese take to be clever responses to their questions. Now, Searle does not understand Chinese. He doesn't know the meaning of any Chinese character. To him, the characters of written Chinese are just a bunch of elaborate squiggles. However, Searle has with him in the room baskets of Chinese characters, and he has a rulebook which says that if he gets certain sequences of characters he should create certain other sequences of characters and slide them out the slot.

Searle intends his thought experiment as a demonstration that syntax is not sufficient to produce semantics. According to Searle,

the room appears to behave as though it understands Chinese; yet neither he nor anything in the room can be said to understand Chinese. There are many arguments for and against Searle's claims, and I will not review them here. Instead, I want to interpret the Chinese room in a completely different way: The Chinese room is a sociocultural cognitive system. The really nice thing about it is that it shows us very clearly that the cognitive properties of the person in the room are not the same as the cognitive properties of the room as a whole. There is John Searle with a basket of Chinese characters and a rulebook. Together he and the characters and rulebook in interaction seem to speak Chinese. But Searle himself speaks not a word of Chinese.

Let us be clear, then, on the distinction between the cognitive properties of the sociocultural system and the cognitive properties of a person who is manipulating the elements of that system.

The heart of Turing's great discovery was that the embodied actions of the mathematician and the world in which the mathematician acted could be idealized and abstracted in such a way that the mathematician could be eliminated. What remained was the essence of the application of rules to strings of symbols. For the purposes of producing the computation, the way the mathematician actually interacted with the material world is no more than an implementational detail. Pylyshyn (1989) claims that while Turing was developing the notion of the mechanically "effective procedure" he was looking "at what a mathematician does in the course of solving mathematical problems and distilling this process to its essentials." The question of what constitutes the essentials here is critical. For Turing, the essentials evidently involve the patterns of manipulations of the symbols, but they expressly do *not* involve the psychological processes which the mathematician uses in order to accomplish the manipulations. The essentials of the abstract manipulation of symbols are precisely not what the person does. What Turing modeled was the computational properties of a sociocultural system.

When the manipulation of symbols is automated, neither the cognitive processes nor the activity of the person who manipulated the symbols is modeled. The symbols themselves are dematerialized and placed inside the machine, or fed to it in a form that permits the straightforward generation of internal representations. What is important about this is that all the problems the mathematician faced when interacting with a world of material symbol to-

kens are avoided. That is good news, if those things are considered unimportant, because they are a nuisance to model anyway. The rulebook (or the mathematician's scribbled notations of rules) is replaced by abstract rules, also inside the computer. The mathematician who was a person interacting with a material world is neither modeled by this system nor replaced in it by something else. The person is simply absent from the system that performs automatic symbol manipulation. What is modeled is the abstract computation achieved by the manipulation of the symbols.

All that is fine if your goal is to extend the boundaries of human computational accomplishments. But it is not necessarily a model of the processes engaged in by a person doing that task. These programs produce the properties, not of the person, but of the sociocultural system. This is a nontrivial accomplishment. But the culture of cognitive science has forgotten these aspects of its past. Its creation myths do not include this sort of analysis. *The physical-symbol-system architecture is not a model of individual cognition. It is a model of the operation of a sociocultural system from which the human actor has been removed.*

Having failed to notice that the central metaphor of the physical-symbol-system hypothesis captured the properties of a sociocultural system rather than those of an individual mind, AI and information-processing psychology proposed some radical conceptual surgery for the modeled human. The brain was removed and replaced with a computer. The surgery was a success. However, there was an apparently unintended side effect: the hands, the eyes, the ears, the nose, the mouth, and the emotions all fell away when the brain was replaced by a computer.

The computer was not made in the image of the person. *The computer was made in the image of the formal manipulations of abstract symbols. And the last 30 years of cognitive science can be seen as attempts to remake the person in the image of the computer.*

It is no accident that the language of the physical-symbol-system hypothesis captures so much of what is happening in domains like ship navigation. The physical-symbol-system hypothesis is based on the operation of systems of this type. Conversely, there is nothing metaphorical about talking about the bearing record book as a memory, or about viewing the erasure of lines drawn in pencil on a chart as forgetting. Sometimes my colleagues ask me whether I feel safe metaphorically extending the language of what's happening

inside people's heads to these worlds. My response is "It's not a metaphorical extension at all." The computer was made in the image of the sociocultural system, and the human was remade in the image of the computer, so the language we use for mental events is the language that we should have used for these kinds of sociocultural systems to begin with. These are not examples of metaphorical extension from the base domain of mental events to the target domain of cultural activity. Rather, the *original* source domain for the language of thought was a particular highly elaborated and culturally specific world of human activity: that of formal symbol systems.

At first, the falling away of the apparatus that connects the person to the world went unnoticed. This may have been because there was a lot of justifiable excitement about what could be done with this technology. All that remained of the person, however, was the boundary between the inside and the outside. And this boundary was not the same as the boundary of the Chinese room. The boundary that remained was assumed to be the boundary of the person—the skin or the skull. In fact, it was the boundary of the formal system. The boundary between inside and outside became the boundary between abstract symbols and the world of phenomena described by symbols. The walls of the Chinese room were mistaken for the skin of the person. And the walls of the room surrounded the symbols, so the symbols were assumed to be inside the head.

This separation between the boundaries of the formal system and the skin shows up in the language of cognitive science. "Symbol systems are an interior milieu, *protected from the external world*, in which information processing in the service of the organism can proceed." (Newell et al. 1989: 107 [my emphasis]). Or:

*Act*, as is typical of many theories of cognition, focuses on the central architecture. Perception and motor behavior are assumed to take place in additional processing systems off stage. Input arrives in working (memory), which thus acts as a buffer between the unpredictable stream of environmental events and the cognitive system. (ibid.: 117)*

The "off stage" metaphor of Newell et al. expresses the isolation of the cognitive system from even the sensory and motor experiences of an organism. In fact, many cognitive scientists take the word 'cognitive' as an antonym to 'perceptual' or 'motor'. Here is a

typical example of this usage: "This is especially true for tasks that are primarily cognitive, in which perceptual and motor operations play only a small role in the total sequence." A strong claim about the modularity of the human cognitive system is implicit in this use of language. It places a large divide between cognition and the world of experience. But the existence of perceptual and motor processes that are distinct and separate from so-called cognitive processes is not an empirical fact: it is simply a hypothesis that was made necessary by having constructed cognition out of a mechanized formal symbol processing system. Proponents of the physical-symbol-system hypothesis point to the existence of various sensory and motor memories that can act as buffers between cognition and the world of experience as evidence of this modularity. In fact, there may be many other uses for such buffers. We are unlikely to discover these other uses, however, as long as we keep cognition isolated from the world. For example, such buffers may be essential in maintaining training signals after the disappearance of stimuli while learning is taking place.

The model of human intelligence as abstract symbol manipulation and the substitution of a mechanized formal symbol-manipulation system for the brain result in the widespread notion in contemporary cognitive science that symbols are inside the head. The alternative history I offer is not really an account of how symbols got inside the head; it is a historical account of how cognitive science put symbols inside the head. And while I believe that people do process symbols (even ones that have internal representations), I believe that it was a mistake to put symbols inside in this particular way. The mistake was to take a virtual machine enacted in the interactions of real persons with a material world and make that the architecture of cognition.

This mistake has consequences. Why did all the sensorimotor apparatus fall off the person when the computer replaced the brain? It fell off because the computer was never a model of the person to begin with. Remember that the symbols were outside, and the apparatus that fell off is exactly the apparatus that supported interaction with those symbols. When the symbols were put inside, there was no need for eyes, ears, or hands. Those are for manipulating objects, and the symbols have ceased to be material and have become entirely abstract and ideational. The notion of abstractness was necessary to bleach the material aspect out of the symbols so that they could become freed from any particular

material instantiation. Calling logic and mathematics "abstract" more than misses the point of their concrete nature as human activities; it obscures it in a way that allows them to be imported into a cognitive inner sanctum. The physicality of material symbols in the environment has been replaced by the physicality (causality) of the computer; thus, while the physical is acknowledged in the physical-symbol-system hypothesis, it is rendered irrelevant by the claim that the physical aspect is an implementational detail. This idea may also help to explain the indifference that cognitive science generally shows to attempts to study implementation in real human systems.

Observe how a proponent of the classical view treats the manipulation of a computational artifact. Here Pylyshyn (1989: 56) constructs an example of manipulations of symbols that are codes for numbers:

*If you can arrange for the computer to transform them systematically in the appropriate way, the transformations can correspond to useful mathematical operations such as addition or multiplication. Consider an abacus. Patterns of beads represent numbers. People learn rules for transforming these patterns of beads in such a way that the semantic interpretation of before-and-after pairs corresponds to a useful mathematical function. But there is nothing intrinsically mathematical about the rules themselves; they are just rules for moving beads around. What makes the rules useful for doing mathematics is that we are assured of a certain continuing correspondence between the formal or syntactic patterns of beads and mathematical objects (such as numbers).*

There are no hands or eyes in this description. There are only the formal properties of the patterns of beads. Pylyshyn is using the example of the abacus to show how the manipulation of symbols produces computations. He provides a very nice illustration of the power of this cultural artifact. He is not interested in what the person does, or in what it means for a person to learn, to "know," or to apply a rule. Rather, he is interested in the properties of the system enacted by the person manipulating the physical beads. That is fine as a description of the computational properties of a sociocultural system, but to take this as being about cognitive processes inside the skin is to recapitulate the error of mistaking the properties of the sociocultural system for the properties of a person. It is easy to do. It is something we do in our folk psychology all the time. But

when one is really careful about talking about cognition, one must carefully distinguish between the tasks that the person faces in the manipulation of symbolic tokens and the tasks that are accomplished by the manipulation of the symbolic tokens.

A failure to do this has led to a biased view of the tasks that are properly considered cognitive. Problem solving by heuristic search is taken as a representative cognitive activity. This is tailor-made for the symbol-shuffling apparatus. The definition of cognition has been unhooked from interaction with the world. Research on games and puzzles has produced some interesting insights, but the results may be of limited generality. The tasks typically chosen for laboratory studies are novel ones that are regarded by subjects as challenging or difficult. D'Andrade (1989) has likened the typical laboratory cognitive tasks to feats of athletic prowess. If we want to know about walking, studying people jumping as high as they can may not be the best approach. Such tasks are unrepresentative in another sense as well. The evolution of the material means of thought is an important component of culturally elaborated tasks. It permits a task that would otherwise be difficult to be re-coded and re-represented in a form in which it is easy to see the answer. This sort of development of material means is intentionally prohibited in puzzle tasks because to allow this sort of evolution would destroy the puzzling aspects of the puzzle. Puzzles are tasks that are preserved in the culture because they are challenging. If the performance mattered, we would learn to re-represent them in a way that removed the challenge. That would also remove their value as puzzles, of course. The point is that the tasks that are "typical" in laboratory studies of thought are drawn from a special category of cultural materials that have been isolated from the cognitive processes of the larger cultural system. This makes these tasks especially unrepresentative of human cognition.

Howard Gardner (1985) is very kind to cognitive science when he says that emotion, context, culture and history were deemphasized in early cognitive science because, although everyone believed they were important, everyone also knew that they complicated things enormously. According to Gardner, getting the program started required a simple model of cognition. The field therefore deferred consideration of affect, culture, context, and history until such time as there was a good model of how an individual worked in isolation. It was hoped that these things could be added in later. That is a charitable reading of the history, I think. I can see why

there were compelling reasons to see it as it was seen, and not to notice that something is wrong when AI was producing "deaf, dumb, and blind, paraplegic agents" (Bobrow 1991) as models of human cognition.

Newell et al. (1989) seemed genuinely puzzled by the fact that no one had succeeded in integrating emotion into the system of cognition they had built. Yet this failure is completely predictable from the assumptions that underlie the construction of the symbol-manipulation model of cognition. The person was simply omitted from what was taken as the model of the cognitive system. The model of cognition came from exactly that part of the system that was material rather than human. Within this underlying theory of cognition there can be no integration of emotion, because the part of the cultural system that is the basis of the physical symbol system excludes emotion. The integration of cognition with action will remain difficult because the central hypothesis separates cognition and action by definition. History and context and culture will always be seen as add-ons to the system, rather than as integral parts of the cognitive process, because they are by definition outside the boundaries of the cognitive system.

Adherents of the physical-symbol-system hypothesis are obviously aware of the presence of a world in which action takes place, and they have attempted to take it into account. Consider the following passage from Newell and Simon's seminal 1972 book *Human Problem Solving* :

*For our theory, specification of the external memories available to the problem solver is absolutely essential. These memories must be specified in the same terms as those we have used for the internal memories; symbol capacities, accessing characteristics, and read and write times. The problem solving program adopted by the information-processing system will depend on the nature of its "built in" internal STM and LTM [short-term memory and long-term memory].*

*From a functional viewpoint, the STM should be defined, not as an internal memory, but as the combination of (1) the internal STM and (2) the part of the visual display that is in the subject's foveal view. . . .*

*In short, although we have few independent data suited to defining precisely how external memory can augment STM, the two components do appear to form a single functional unit as far as the*

*detailed specification of a problem solving information-processing system is concerned.*

This is a good start on the problem, but I think it is fair to say that in the twenty years since the publication of *Human Problem Solving* the use of material structure in the problem-solving environment has not been a central topic in the physical-symbol-system research agenda. Some recent work within this tradition takes the "external world" into account (Larkin 1989; Vera and Simon 1993) but treats the world only as an extra memory on which the same sorts of operations are applied as are applied to internal memories. Structure in the world can be much more than an augment to memory. The use of cultural structures often involves, not just the same process with more memory, but altogether different processes. The overattribution of internal structure results from overlooking the coordination of what is inside with what is outside. The problem remains that the nature of the interaction with the world proposed in these systems is determined by the assumptions of the symbolic architecture that require the bridging of some gap between the inner, cognitive world and an outer world of perception and action.

These criticisms by themselves are not sufficient grounds for rejecting the notion that humans are symbol-processing systems. Newell and Simon (1990) wisely acknowledge that the physical-symbol-system hypothesis is a hypothesis and that the role of symbolic processes in cognition is an empirical question. It has proved possible to interpret much of human problem-solving behavior as if the very architecture of human cognition is symbol processing. It's a hypothesis. A lot of data can be read as failing to reject it. Yet, the hypothesis got there under suspicious circumstances. There are no plausible biological or developmental stories telling how the architecture of cognition became symbolic. We must distinguish between the proposition that the architecture of cognition is symbolic and the proposition that humans are processors of symbolic structures. The latter is indisputable, the former is not. I would like to be able to show how we got to be symbol manipulators in relation to how we work as participants in sociocultural systems, rather than assume it as an act of faith. The origins of symbolic processes have not been explored this way, though, because they were obfuscated by the creation myth that maintains that the computer was made in the image of humans.

Increasingly, the physical-symbol-system hypothesis is a perspective into which things don't fit. It was a bet or a guess, grounded in a nearly religious belief in the Platonic status of mathematics and formal systems as eternal verities rather than as historical products of human activity. This is an old dispute that lies at the heart of the developing split in cognitive science between those who feel there is more to be learned from the physical-symbol-system research and those who feel it has been exhausted. (See the special January-March 1993 issue of *Cognitive Science*.) By advocating this alternative view, I am espousing what might be called a "secular" view of cognition—one that is grounded in a secular perspective on formal systems, in contrast with the quasi-religions "cosmic truth" view put forth by the symbolists.

Why cognition became disembodied is clear from the history of the symbolic movement. An important component of the solution is to re-embody cognition, including the cognition of symbol processing.

I believe that humans actually process internal representations of symbols. But I don't believe that symbol manipulation is the architecture of cognition. Historically, we simply assumed that symbol processing was inside because we took the computer as our model of mentality. Humans (and, I suspect, most other animals) are good at detecting regularities in their environment and at constructing internal processes that can coordinate with those regularities. Humans, more than any other species, spend their time producing symbolic structure for one another. We are very good at coordinating with the regularities in the patterns of symbolic structure that we present to one another. As was described in chapter 7, the internal structures that form as a consequence of interaction with symbolic materials can be treated as symbolic representations. Ontogenetically speaking, it seems that symbols are in the world first, and only later in the head.

## Studying Cognition in the Wild

Many of the foundational problems in cognitive science are consequences of our ignorance of the nature of cognition in the wild. Most of what we know about cognition was learned in laboratory experiments. Certainly, there are many things that can be learned only in closely controlled experiments. But little is known about the relationships of cognition in the captivity of the laboratory to

cognition in other kinds of culturally constituted settings. The first part of the job is, therefore, a descriptive enterprise. I call this description of the cognitive task world a "cognitive ethnography." One might have assumed that cognitive anthropology would have made this sort of work its centerpiece. It has not. Studying cognition in the wild is difficult, and the outcomes are uncertain.

Cognitive systems like the one documented in this book exist in all facets of our lives. Unfortunately, few truly ethnographic studies of cognition in the wild have been performed. (Beach 1988, Frake 1985, Gladwin 1970, Goodwin 1993, Goodwin and Goodwin 1992 and 1995, Latour 1986, Lave 1988, Lave et al. 1984, Ochs et al. (in press), Scribner 1984, Suchman 1987, and Theureau 1990 are lonely exceptions to this trend.) We trust our lives to systems of this sort every day, yet this class of phenomena has somehow fallen into the crack between the established disciplines of anthropology and psychology and appears to be excluded by foundational assumptions in cognitive science. This book is an attempt to show what a natural history of a cognitive system could be like.

Among the benefits of cognitive ethnography for cognitive science is the refinement of a functional specification for the human cognitive system. What is a mind for? How confident are we that our intuitions about the cognitive nature of tasks we do on a daily basis are correct? It is a common piece of common sense that we know what those tasks are because we are human and because we engage in them daily. But I believe this is not true. In spite of the fact that we engage in cognitive activities every day, our folk and professional models of cognitive performance do not match what appears when cognition in the wild is examined carefully. I have tried to show here that the study of cognition in the wild may reveal a different sort of task world that permits a different conception of what people do with their minds.

Cognitive science was born in a reaction against behaviorism. Behaviorism had made the claim that internal mental structure was either irrelevant or nonexistent—that the study of behavior could be conducted entirely in an objective characterization of behavior itself. Cognitive science's reaction was not simply to argue that the internal mental world was important too; it took as its domain of study the internal mental environment largely separated from the external world. Interaction with the world was reduced to read and write operations conducted at either end of extensive processing

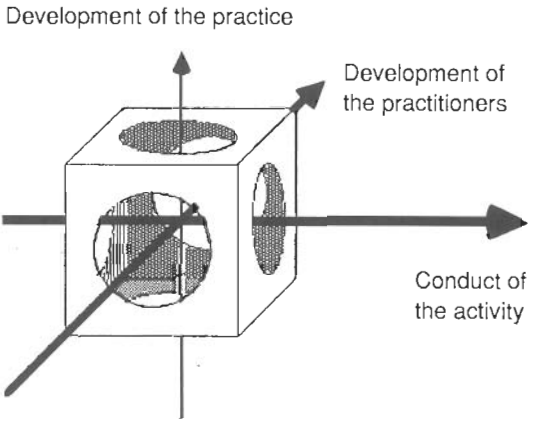


**Figure 9.1** A moment of human practice.

activity. This fit the computer metaphor very well, but it made the organization of the environment in which thinking took place seem largely irrelevant. Both behaviorism and cognitivism must be wrong.

## Cognition in the Intersection of Cultural Processes

The cube depicted in figure 9.1 represents any moment in navigation practice (or, in fact, any moment in any human practice). The arrows passing through the cube represent three developmental sequences of which every moment of practice is simultaneously a part. I have adopted some simple conventions to capture several aspects of the situation in this single diagram. The thickness of the arrow represents the density of interaction among the elements in that dimension. The length of the shaft of the arrow emerging from the cube represents the rate at which states in that dimension are changing. The length of the tail of the arrow going into the cube represents the duration of the relevant history of the activity in the given dimension.

It is essential to keep in mind that these things are all happening at the same time in the same activity. Having reinstated a whole human being in a culturally constituted activity, I see the following.

The conduct of the activity proceeds by the operation of functional systems that bring representational media into coordination with one another. The representational media may be inside as

well as outside the individuals involved. These functional systems propagate representational state across the media. In describing the ongoing conduct of navigation tasks, it is possible to identify a number of cognitive systems, some subsuming others. One may focus on the processes internal to a single individual, on an individual in coordination with a set of tools (chapter 3), or on a group of individuals in interaction with one another and with a set of tools (chapter 4). Each system produces identifiable cognitive properties, and in each case the properties of the system are explained by reference to processes that transform states inside the system (chapter 5). The structured representational media in the system interact in the conduct of the activity. Each medium is put to use in an operational environment constituted by other media. As indicated in figure 9.1, the conduct of the activity itself has a relatively short history. An entry into a harbor, for example, involves a few hours of preparation and takes about an hour to complete. Changes in this dimension happen quickly, and the elements of the task performance are in relatively intense interaction with one another. The conduct of the activity creates elements of representational structure that survive beyond the end of the task. These elements—the logbooks, pencil marks on charts, the quartermasters' memories of the events—are the operational residua of the process.

In this adaptive system, the media may be changed by the very processes that constitute the conduct of the activity. The operations of the navigation team produce a structured experience for the participants that contains opportunities for individual learning (chapter 6). As a consequence of their participation in the task performance, the quartermasters may acquire internal organization that permits them to coordinate with the structure of their surroundings. In this way, learning can be seen as the propagation of organization through an adaptive system (chapter 7). The development of the practitioners themselves takes years. Through a career, a quartermaster gradually acquires the skills that are exercised in the performance of the job. Changes to the organization of the internal media that the quartermasters bring to the job take place more slowly than the changes to the states that the media support. That is, it takes longer to learn how to plot a fix, for example, than it does to plot a fix. But since most learning in this setting happens in the doing, the changes to internal media that permit them to be coordinated with external media happen in the same processes that

bring the media into coordination with one another. The changes to the quartermasters' skills and the knowledge produced by this process are the mental residua of the process.

The setting of navigation work evolves over time as partial solutions to frequently encountered problems are crystallized and saved in the material and conceptual tools of the trade and in the the social organization of the work. The development of the practice takes place over centuries (chapter 2). The very same processes that constitute the conduct of the activity and that produce changes in the individual practitioners of navigation also produce changes in the social, material, and conceptual aspects of the setting. The example given in chapter 8 illustrates the creation in interaction of a new concept and a shared lexical label for it (the "total" in the modular form of the true-bearing computation). The microgenesis of the cultural elements that make up the navigation setting is visible in the details of the ongoing practice.

All this happens simultaneously in cognition in the wild. It is in this sense that cognition is a fundamentally cultural process.