

This excerpt from

Cognition in the Wild.
Edwin Hutchins.
© 1995 The MIT Press.

is provided in screen-viewable form for personal use only by members of MIT CogNet.

Unauthorized use or dissemination of this information is expressly forbidden.

If you have any questions about this material, please contact cognetadmin@cognet.mit.edu.

Introduction

The seed from which this book grew was planted in November 1980, when I spent most of a day on the navigation bridge of a U.S. Navy ship as it worked its way in from the open North Pacific, through the Straits of Juan de Fuca, and down Puget Sound to Seattle. I was aboard the ship to study what the operators of its steam propulsion plant knew and how they went about knowing it. I had spent most of the preceding week down in the bowels of the ship, observing engineering operations and talking to the boiler technicians and machinist's mates who inhabited that hot, wet, noisy tangle of boilers, pumps, and pipes called the engineering spaces. I'll admit to having felt a little claustrophobic after all that time spent below the water line, where there is no night or day and the only evidence of being at sea is the rhythmic tipping of the deck plates and sloshing of water in the bilge below one's feet as the ship rolls in the swell. A chief boiler technician confided to me that in 21 years on Navy ships he had never yet been on deck to experience either of those two most romantic seafaring events, a ship's arrival at or departure from a port.

I resolved, therefore, to take my last few hours aboard this ship on the navigation bridge, where I could see out the windows or even go out on the bridge wing to get a breath of cold fresh air. My professional rationalization for being on the bridge was that there I would be able to observe the process that generates the flurry of engine commands that always taxes the engineering crew when the ship nears the dock. And I did make a detailed record of all engine and helm commands given in the 75 minutes from the time the engines were first slowed until they were secured—there were 61 in all. But what really captured my attention was the work of the navigation team.

Three and a half years later, the project that became this book began in earnest. In the summer of 1984, I was still working for the Navy Personnel Research and Development Center in San Diego as a civilian scientist with the title Personnel Research Psychologist. By then I had participated in two successful and well-known

projects. With these successes came the freedom to conduct an independent research project. I was given *carte blanche* to study whatever I thought was of most interest. I chose to study what I was then calling *naturally situated cognition*. Having a research position in a Navy laboratory made it possible for me to gain access to naval vessels, and my longtime love of navigation and experience as a racing yacht navigator made it easy for me to choose navigation as an activity to study afloat. I talked my way aboard a ship and set up shop on the navigation bridge. At the time, I really had no notion what an ideal subject navigation would turn out to be. When I began, I was thinking in terms of the naturally situated cognition of individuals. It was only after I completed my first study period at sea that I realized the importance of the fact that cognition was socially distributed.

A little earlier, I had been asked to write a book describing what is in cognitive anthropology for the rest of cognitive science. I began that project, but after I became disillusioned with my field I lost interest in it. The choice of naturally situated cognition as a topic came from my sense that it is what cognitive anthropology really should have been about but largely had not been. Clifford Geertz (1983) called for an “outdoor psychology,” but cognitive anthropology was unable or unwilling to be that. The respondents may have been exotic, but the methods of investigation were largely borrowed from the indoor techniques of psychology and linguistics. When cognitive and symbolic anthropology split off from social anthropology, in the mid 1950s, they left society and practice behind.

As part of the cognitive revolution, cognitive anthropology made two crucial steps. First, it turned away from society by looking inward to the knowledge an individual had to have to function as a member of the culture. The question became “What does a *person* have to know?” The locus of knowledge was assumed to be inside the individual. The methods of research then available encouraged the analysis of language. But knowledge expressed or expressible in language tends to be declarative knowledge. It is what people can say about what they know. Skill went out the window of the “white room.” The second turn was away from practice. In the quest to learn what people know, anthropologists lost track both of how people go about knowing what they know and of the contribution of the environments in which the knowing is accomplished. Perhaps these narrowing assumptions were necessary to

get the project of cognitive anthropology off the ground. I will argue that, now that we are underway as a discipline, we should revoke these assumptions. They have become a burden, and they prevent us from seeing the nature of human cognition.

In particular, the ideational definition of culture prevents us seeing that systems of socially distributed cognition may have interesting cognitive properties of their own. In the history of anthropology, there is scarcely a more important concept than the division of labor. In terms of the energy budget of a human group and the efficiency with which a group exploits its physical environment, social organizational factors often produce group properties that differ considerably from the properties of individuals. Clearly, the same sorts of phenomena occur in the cognitive domain. Depending on their organization, groups must have cognitive properties that are not predictable from a knowledge of the properties of the individuals in the group. The emphasis on finding and describing “knowledge structures” that are somewhere “inside” the individual encourages us to overlook the fact that human cognition is always situated in a complex sociocultural world and cannot be unaffected by it.

Similar developments in the other behavioral sciences during the cognitive revolution of the late 1950s and the 1960s left a troubled legacy in cognitive science. It is notoriously difficult to generalize laboratory findings to real-world situations. The relationship between cognition seen as a solitary mental activity and cognition seen as an activity undertaken in social settings using various kinds of tools is not at all clear.

This book is about softening some boundaries that have been made rigid by previous approaches. It is about locating cognitive activity in context, where context is not a fixed set of surrounding conditions but a wider dynamical process of which the cognition of an individual is only a part. The boundaries to be softened or dissolved have been erected, primarily for analytic convenience, in social space, in physical space, and in time. Just as the construction of these boundaries was driven by a particular theoretical perspective, their dissolution or softening is driven by a different perspective—one that arose of necessity when cognition was confronted in the wild.

The phrase “cognition in the wild” refers to human cognition in its natural habitat—that is, to naturally occurring culturally constituted human activity. I do not intend “cognition in the wild” to

be read as similar to Lévi-Strauss's "pensée sauvage," nor do I intend it to contrast with Jack Goody's (1977) notion of domesticated mind. Instead, I have in mind the distinction between the laboratory, where cognition is studied in captivity, and the everyday world, where human cognition adapts to its natural surroundings. I hope to evoke with this metaphor a sense of an ecology of thinking in which human cognition interacts with an environment rich in organizing resources.

The attempt is cultural in nature, giving recognition to the fact that human cognition differs from the cognition of all other animals primarily because it is intrinsically a cultural phenomenon. My aim is to provide better answers to questions like these: What do people use their cognitive abilities for? What kinds of tasks do they confront in the everyday world? Where shall we look for explanations of human cognitive accomplishment?

There is a common misconception among cognitive scientists, especially those who do their work in laboratory settings, that research conducted outside the laboratory is necessarily "applied" work. I will argue in what follows that there are many excellent reasons to look at the "real world" that are not concerned with hoped-for applications of the research findings (although funding sponsors often like to think in those terms). Pure research on the nature of real cognitive practices is needed. In this book, I emphasize practice not in order to support a utilitarian or functionalist perspective but because it is in real practice that culture is produced and reproduced. In practice we see the connection between history and the future and between cultural structure and social structure. One of my goals in writing this book is to make clear that the findings of pure research on cognition in the wild should change our ideas about the nature of human cognition in general. This is not news to anthropologists, who have been doing pure research in the form of ethnography for decades.

This book is an attempt to put cognition back into the social and cultural world. In doing this I hope to show that human cognition is not just influenced by culture and society, but that it is in a very fundamental sense a cultural and social process. To do this I will move the boundaries of the cognitive unit of analysis out beyond the skin of the individual person and treat the navigation team as a cognitive and computational system.

Chapter 1, "Welcome Aboard," attempts to locate the activity of ship navigation in the larger world of modern life. It weaves to-

gether three journeys: a movement through physical space from the “street” to the ship, a movement through social space from civilian to military life, and a movement through conceptual space from everyday notions of wayfinding to the technical domain of navigation. Both the researcher and the reader must make these journeys to arrive at the activity of navigation as practiced on the bridge of a Navy ship. Military ranks and the ways in which military identities are formed are presented here because these things affect individual’s relationships to their work. An important aspect of the larger unit is that it contains computational elements (persons) who cannot be described entirely in computational terms. Who they talk to and how they talk to one another depend on these social organizational factors. This chapter also contains a discussion of the relationship of the researcher to the activity under study. (The name of the ship and the names of all the individuals mentioned in the book are pseudonyms. All the discourses reported, whether standing alone in transcript form or embedded in narrative passages were transcribed directly from audio recordings of actual events.)

Having taken navigation as it is performed by a team on the bridge of a ship as the unit of cognitive analysis, I attempt in chapter 2, “Navigation as Computation,” to apply the principal metaphor of cognitive science—cognition as computation—to the operation of this system. I should note here that in doing so, I do not make any special commitment to the nature of the computations that are going on inside individuals except to say that whatever happens there is part of a larger computational system. This chapter describes the application of David Marr’s notions of levels of analysis of cognitive systems to the navigation task and shows that, at the computational level, it is possible to give a single description of the computational constraints of all known technical forms of human navigation. A comparison of modern Western navigation with navigation as practiced in Micronesia shows that considerable differences between these traditions lie at the representational/algorithmic level and at the implementational level. A brief historical review of the development of modern navigation shows that the representational and implementational details of contemporary practice are contingent on complex historical processes and that the accumulation of structure in the tools of the trade is itself a cognitive process.

Chapters 3–5 explore the computational and cognitive properties of systems that are larger than an individual. The issues addressed

in these chapters concern how these larger systems operate and how their cognitive properties are produced by interactions among their parts.

Chapter 3, "The Implementation of Contemporary Pilotage," describes the physical structures in which the navigation computations are implemented. This chapter elaborates a conception of computation as the propagation of representational state across a variety of media. This view of computation permits the use of a single language of description to cover cognitive and computational processes that lie inside and outside the heads of the practitioners of navigation. The first section of this chapter describes the "fix cycle" as a cognitive process. The second section describes how navigation tools are used and how local functional systems composed of a person in interaction with a tool have cognitive properties that are radically different from the cognitive properties of the person alone. The third section discusses the ways in which the computational activity can be distributed through time by pre-computing not only partial results but also the means of computation. I show here how the environments of human thinking are not "natural" environments. They are artificial through and through. Humans create their cognitive powers by creating the environments in which they exercise those powers. This chapter concludes with a discussion of the relationship between the cognitive properties of the individuals performing a task and the cognitive properties of the system in which they participate.

Chapter 4, "The Organization of Team Performances," moves the boundaries of the unit of analysis even further out to consider the cognitive properties of the team as a whole. Here I note some of the problems that are encountered when cognitive activities are distributed across the members of a group. It is not the case that two or more heads are always better than one. This chapter describes the structures and processes involved in the group performance of the navigation task. The first section follows through on the application of Marr's concepts of computation to the navigation activity and discusses the properties of the activity as an explicitly computational system. The second section presents a problem in work organization encountered by the navigation team and shows why it is often difficult to apply the concepts that organize individual action to the organization of group action. The final section shows how the members of the navigation team form a flexible connective tissue that maintains the propagation of representational state in the face of a range of potentially disruptive events.

Chapter 5, “Communication,” continues the theme of chapter 4 but looks at communication in more detail. It asks: How is it that patterns of communication could produce particular cognitive properties in a group? The chapter begins with a discussion of features of communication observed in the navigation team and their effects on the Team’s computational properties. These observations lead to some simple hypotheses about the ways in which patterns of communication might affect the computational properties of a group. These hypotheses are explored using a computer simulation of communities of connectionist networks. The simulations lead to the surprising conclusion that more communication is not always better.

Chapters 6–8 concern learning or change in the organization of cognitive systems at several scales.

Chapter 6, “The Context of Learning,” is a bridge between the descriptions of ongoing operations provided by the previous chapters and the descriptions of changes in the nature of ongoing operations provided by the following chapters. It describes the context in which novice navigators become experts. This chapter is an attempt to examine both the work that the system does in order to scaffold learning by practitioners and the opportunities for the development of new knowledge in the context of practice.

Whereas in chapter 6 I deal with the observable contexts surrounding learning, in chapter 7, “Learning in Context,” I try to dissolve the boundaries of the skin and present navigation work as a system of interactions among media both inside and outside the individual. I look at learning or conceptual change as a kind of adaptation in a larger dynamical system. This chapter presents a functional notation and a framework for thinking about learning as local adaptation in a dynamic system of coordinations of representational media.

Chapter 8, “Organizational Learning,” returns the focus to the larger unit of analysis: the team as a whole. It presents a case study of an incident in which the navigation team was forced to adapt to changes in its information environment. The analysis presented here examines a particular incident in which the microstructure of the development of the navigation practice can be seen clearly. It is an attempt to show the details of the kinds of processes that must be the engines of cultural change.

Chapter 9, “Cultural Cognition,” attempts to pull the preceding chapters together into a coherent argument about the relationships of culture and cognition as they occur in the wild. I attempt first

to illustrate the costs of ignoring the cultural nature of cognition. I argue that a new framework is needed to understand what is most characteristically human about human cognition. In order to construct a new framework, the old one must be deconstructed. I therefore provide two readings of the history of cognitive science: a history as seen by the proponents of the currently dominant paradigm and a rereading of the history of cognitive science from a sociocultural perspective. The differences between these two readings highlight a number of problems in contemporary cognitive science and give new meanings to some of the familiar events in its history.

This excerpt from

Cognition in the Wild.
Edwin Hutchins.
© 1995 The MIT Press.

is provided in screen-viewable form for personal use only by members of MIT CogNet.

Unauthorized use or dissemination of this information is expressly forbidden.

If you have any questions about this material, please contact cognetadmin@cognet.mit.edu.