Interfacing Thought

Cognitive Aspects of Human-Computer Interaction

edited by John M. Carroll

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

This book was set in Palatino by Asco Trade Typesetting Ltd., Hong Kong, and printed and bound by Halliday Lithograph in the United States of America.

Library of Congress Cataloging-in-Publication Data

Interfacing thought.
"A Bradford book."
Includes bibliographies and index.
1. Interactive computer systems. 2. Human engineering. I. Carroll, John M., 1950–.
QA76.9.I58I59 1987 004'.01'9 86-20974
ISBN 0-262-03125-6

Interfacing Thought

Contributors

John R. Anderson
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA

Philip J. Barnard
MRC Applied Psychology Unit
Cambridge, UK

John B. Black
Teachers College
Columbia University
New York, NY

John M. Carroll
IBM Research
Yorktown Heights, NY

Mitchell DeRidder
Department of Psychology
University of Texas
Austin, TX

Donald J. Foss
Department of Psychology
University of Texas
Austin, TX

Dana S. Kay
Computer Science Department
Yale University
New Haven, CT

Thomas K. Landauer
Bell Communication Research
Morristown, NJ

Clayton H. Lewis
Computer Science Department
University of Colorado
Boulder, CO

Jean McKendree
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA

Thomas W. Malone
Sloan School of Management
Massachusetts Institute of Technology
Cambridge, MA

Richard E. Mayer
Department of Psychology
University of California
Santa Barbara, CA

Donald A. Norman
Institute for Cognitive Studies
University of California
San Diego, CA

Judith Reitman Olson
Graduate School of Business
Administration
University of Michigan
Ann Arbor, MI

Peter G. Polson
CLIPR
University of Colorado
Boulder, CO

Phyllis Reisner
IBM Almaden Research Center
San Jose, CA

Mary Beth Rosson
IBM Research
Yorktown Heights, NY

Elliot M. Soloway
Computer Science Department
Yale University
New Haven, CT

John Whiteside
Digital Equipment Corporation
Nashua, NH

Dennis Wixon
Digital Equipment Corporation
Nashua, NH

Preface

John M. Carroll

One of the most active new research areas in cognitive science is that of human-computer interaction. This work concerns itself with the design of "user interfaces," that is, computers as experienced and manipulated by human users. Its principal goal is to provide an understanding of how human motivation, action, and experience place constraints on the usability of computer equipment. The area has grown prodigiously in the last several years.

Computers Are Culturally Pervasive One obvious reason for this growth is that computers are playing increasingly prominent and direct roles in the conduct of human affairs. The emergence of the microcomputer has brought computer presence into our living rooms, kitchens, and game rooms—to say nothing of our workplaces. Computer systems are no longer remote, enormous, and unknown gargantuans to which decks of cards are submitted through intermediaries. They are diverse; they are personal; they talk to us on the phone; they are in our cars. They are an important part of life to a rapidly increasing number of people. All this has made computer problems into important problems.

In the 1930s social scientists concerned themselves with the effects of the automobile on people's lives. In the 1950s and '60s, attention was directed at the effects of television. Computers promise to have a cultural impact at least as great as these.

Computers Can Be Problematic for People The human implications of human-computer interaction include many novel learning and skill-adjustment problems. Indeed, the pervasiveness of computing devices in daily life can be as much a frightening and difficult obstacle as it can be an exciting and empowering development. People often experience serious and frustrating difficulties when they attempt to use computers. Tasks they formerly could accomplish now must be relearned; competent secretaries, accountants, and lawyers are—at least temporarily—returned to varying levels of incompetence until they can master "the system." People often see their difficulties with computers as arbitrarily imposed hurdles that exclude them from participation in activities.

These learning and skill-performance problems are now clearly acknowledged in the industry. Where once new computer systems were developed to store and retrieve more information faster, they are now hawked as "easy to use." Where once the spotlight was on the hardware internals of computers, it has largely shifted to the user interface, to the computer as it is experienced and manipulated by the person using it. This reflects an awareness of the problems people experience, but it is all too often wishful thinking; labeling a computer "easy to use" does not make it easy to use. Clearly these are questions in cognitive science: What are the bases of learning and skill problems? How can these problems be eased?

Integrating Basic Cognitive Theory and Application The chapters in this volume provide an interim report on the project of establishing an applied science of human-computer interaction grounded in the framework of cognitive science. They also explore the intriguing, complementary prospect that applied research in human-computer interaction can contribute in a variety of ways to the development of cognitive science itself. For example, Tom Landauer, in his chapter, observes that basic cognitive psychology has been chronically impeded by the lack of an applied discipline through which to determine the boundary conditions and completeness of its concepts and descriptions. One example he discusses is applying laboratory results on decision time and motor selection to the practical task of designing computer menu structures for ease of use.

Don Norman makes the point that an applied science of human-computer interaction turns out to need all of basic cognitive science and more besides. When we come to basic cognitive theory for principles that can be applied to problems in human-computer interaction, we often find far less than we need: What makes things easier to learn? Which designs accord with natural inclinations of human skill and which do not? It seems like basic cognitive science ought to inform us on these matters, but quite often the basic science that exists is paradigm bound to simple, artificial tasks and cannot be usefully scaled up. Landauer's review of research on names for computer commands also makes this point.

In basic cognitive science it is possible to study an area like word recognition as if human language behavior and experience really occurred word by word, but for an applied endeavor such an unrealistic idealization is unacceptable; it is too obvious that the real situation is more complicated. To make an adequate analysis, we need to take multiple disciplinary perspectives—psychology, linguistics, anthropology, sociology, and computer science—*simultaneously*. As Norman suggests, human-computer interaction, as applied cognitive science, begins to appear as a superset of basic cognitive science!

Developing Cognitive Theory Even within the conventionally recognized scope of cognitive science, theoretical work on human-computer interaction is extending theory. Don Foss and Mitch DeRidder address the metatheoretical problem of choosing between alternative user interface designs as an example of the general problem of choosing between alternative cognitive theories. An interface design after all is a codification of a theory about what will be easy for people to learn and to use. They propose that selection criteria for choosing between alternatives be based on measured transfer of learning; the ease of transition between knowledge states in learning an interface would determine the adequacy of that interface design (and of the cognitive theory it codifies).

Peter Polson's approach to the analysis of transfer of learning is to develop an articulate cognitive description of what is learned (represented as a set of "productions") and then to test specific empirical predictions under the assumption that a production once learned is automatically available in later learning. This formalizes the everyday experience that prior knowledge has an impact on future learning efforts, sometimes facilitating learning and sometimes interfering with learning. Traditional efforts to understand this important phenomenon were often directed at simple and artificial domains (the old standby of nonsense word lists). Polson's work demonstrates how a cognitive description of transfer can be applied to the learning of commands in a computer application, and indeed how such a theory might predict the cognitive consequences of user interface designs before those designs are ever implemented.

The management of limited resources in human performance is another traditional area of theoretical concern in cognitive science. Phil Barnard's chapter argues that cognitive science is in need of a more integrated theory of knowledge representation and information processing—one that focuses more attention on the coordinated control of sequences of behavior in relation to complex and dynamic environments, one that is less paradigm bound with respect to empirical data. Barnard outlines a distributed architecture for cognitive processes called "interacting cognitive subsystems" and applies it to the description of learning computer dialog structures (like menus and command languages). Like Foss and DeRidder and like Polson, Barnard makes the point that current limitations in cognitive theory make the application of theory piecemeal and heuristic. The approach in these three chapters is to develop the general theory base by taking the domain of human-computer interaction as exemplary.

Building Theoretical Perspectives from Established Research Areas Crosscutting the perspective of general cognitive theory is the complementary perspective of particular task domains. One of the important themes in recent cognitive science research is that particular task domains have their own distinctive structures. According to this view, general studies of, say, problem solving need to be augmented by studies of particular domains within which problem solving takes place (a well-known example is the domain of chess). The chapters by Clayton Lewis and by John Black, Dana Kay, and Elliot Soloway examine problems in human-computer interaction using prior cognitive research in other domains as a sort of sounding board. Lewis examines the similarities and the differences between learning mathematics and learning to use a computer application; Black, Kay, and Soloway examine those between reading and understanding a story and understanding computer text editors and programs.

This work is important in that it can help place aspects of the human-computer interaction domain within a larger space of task domains studied in cognitive science. It can help to identify that which is general across various domains and that which is particular to given domains. The principles brought to light in any given investigation have boundary conditions, but these will remain unknown until the principles are studied in a variety of situations. General human strategic biases and the particulars of the computer situation—the tasks, the goals, the available methods, the likely confusions—all interact powerfully and richly in determining how people construe their current situation, what they choose to do, and how they try to do it.

Broadening the Scope of Current Cognitive Science Analyses Several of the chapters in this volume undertake projects extending the scope of current cognitive science research. Tom Malone introduces the concept of "organizational interface" to refer to situations in which groups of people interact by means of computers (e.g., electronic mail applications). The point of organizational interfaces is that people do not perform routine tasks and solve problems by themselves. Traditional examinations of human performance—and of performance with computer systems in particular—has ignored organizational factors. Malone develops the organizational interface concept to suggest ways in which computer systems can be designed to facilitate group problem solving, and conversely to anticipate likely changes in how groups of people will solve problems in the future using computer systems.

Judy Olson extends current analytical techniques by assimilating global task analysis to current cognitive techniques for modeling lower-level aspects of planning. Many theorists have proposed schemes for analyzing human performance in terms of representations of tasks and of human information processing operations. For example, the transfer of learning work discussed by Polson and by Foss and DeRidder models human-computer interaction by representing how a computer application works with respect to how human cognitive capabilities and limitations are engaged. Typically, this work focuses on fairly low levels of task planning and execution (e.g., how to delete a line in an electronic document). Olson incorporates global task analysis, a consideration of tasks at the level of workplace organization, into this framework.

Another example of broadening the scope of current cognitive science analysis is the explicit role that motivation plays in the chapters by Malone and by myself and Mary Beth Rosson. Cognitive tasks like learning and problem solving are frequently analyzed in purely cognitive terms, that is, in terms of information transactions. However, it is clear from everyday experience, and all too salient to researchers in human learning and problem solving, that motivation may play at least as prepotent a role as any purely cognitive factor. A motivated learner often cannot be stopped—even by rather poorly designed user interface facilities. And a poorly motivated learner often cannot be helped. Bringing the consideration of such factors into the science of human-computer interaction from the first broadens the conventional outlook of cognitive science.

Making Ecological Tests of Cognitive Theories Human-computer interaction provides cognitive science with a good test bed for examining the ability of cognitive theories to apply in a complex, real domain. Jean McKendree and John Anderson take this approach in their study of learning to program in Lisp. The programming task is a complex example for applying and developing Anderson's ACT* theory of human cognitive skill. McKendree and Anderson find that speedup in programming tasks consists of two components, general strengthening of previously learned skill components and situation-specific new learning, which consists of organizing new skill components and combining frequently exercised skill components. McKendree and Anderson describe a mathematical model for their learning data, stressing that at a sufficiently fine grain of skill and knowledge analysis, simple and general mechanisms provide a very close fit.

Richard Mayer discusses a series of studies of learning to program in BASIC. Mayer examines the contrasts and relationships between conceptual and syntactic knowledge in the programming domain. He analyzes specific domain-relevant conceptions in terms of their impact on learning and their trainability. This work illustrates how a cognitive approach to learning can provide instructional designs for real training domains.

Guiding the Introduction of New Technology One of the exciting properties of the human-computer interaction domain is its rapid change. As technologies evolve, and as the cultural context vis-à-vis computing evolves in consequence, new research questions are raised, addressed, refined, altered, and finally superseded—as the technological and cultural givens again change. While this is not the sort of field one would enter to study endless parametric variations on some particular laboratory phenomenon, it also is not the sort of field that could ever become static! Because of all this it also represents a rather unusual opportunity for cognitive science: the opportunity to *change* future technology by producing an understanding of contemporary technology—and thereby perhaps to affect the future directly and constructively.

Here lies an important contrast between the direction and ambition of current research on human-computer interaction and earlier studies of the human impact of technological developments (the automobile and television). In these former cases, social scientists were merely observers and analysts. In the case of computer technology, this role is still available, and there are many descriptions of how computers are affecting contemporary life. But there is also a new role to play: Computer technology is more accessible, and more flexible; it is developing faster and with greater diversity. Cognitive scientists can observe and analyze these developments, but they can also participate in and indeed direct these developments.

Yet this may be more even than an opportunity; it is almost a call to arms. For, as stated above, it seems likely that computing technology will have a very great and continuing impact on our world. The problems that arise will be resolved in some way whether or not cognitive scientists provide the methodological and conceptual leverage to address them—there were user interfaces long before there were appreciable numbers of cognitive scientists trying to determine how to enhance usability!

It is of course too early to tell how significant and sustained may be the impact of cognitive science on applied work in human-computer interaction, or what reciprocal effect this applied work might have on cognitive science. Nevertheless, we should not mistake the opportunity and the challenge that this confluence of interests embodies. Mature science provides real empirical leverage in application, and serious applied work is systematically grounded in science. A gauntlet has been thrown.

This book grew out of a symposium, "The Computer: Tool and Topic," that I organized for the 1983 annual meeting of the Society for Computers in Psychology. John Anderson, Tom Landauer, Clayton Lewis, and Don Norman participated as speakers, and several of the other authors in this volume participated with questions from the floor.

Phyllis Reisner and John Whiteside helped to review the chapters, providing technical and editorial feedback that greatly improved them. Phyllis and John also agreed to contribute discussion chapters to facilitate further the kind of technical dialog that is often difficult to achieve in an edited book, but that seems especially important in a new research area.

Finally, I am grateful to my colleagues at the User Interface Institute for keeping me on track in a multitude of ways, and to the IBM Corporation for providing support for this project in the form of my time.

Contents

1 Relations between Cognitive Psychology and Computer System Design

Thomas K. Landauer

Cognitive psychology is more intimately related to the design of computers than to that of traditional machines, such as automobiles and home appliances. There are several reasons. First, the new information technology is so flexible that functions change with bewildering frequency. It is ever less feasible to count on the existence of experienced operators. Unlike typewriters and automobiles, it seems unlikely that information machines of the future will stay the same long enough for public school training to prepare people for lifelong careers based on their use. Thus easy learning or self-evident operations are critical. Second, and equally important, the tasks for which computers are the tools are generally ones in which the human's thought processes themselves are being aided. The maturation of computer applications is taking us ever farther in this direction. The first jobs for computers involved routine information tasks like bookkeeping, in which mechanical procedures once done by humans could simply be assumed by machines. Computers increasingly are used to support dynamic interactive tasks, like text editing and financial simulation, in which the user's mind is an important and lively component of the total system. Designing tools for this kind of activity is an intimately cognitive-psychological activity. Its accomplishment can no longer be viewed as that of first designing a machine to do something, then designing the controls by which the operator guides the machine.

Although the need for greater consideration of users has been recognized for some time now, the response so far, by and large, has been shallow. In attempting to provide greater "user friendliness," designers and programmers have indeed paid more attention to the usability of their systems, and in doing so have exploited the much expanded power of the systems with which they work. For example, they often use the larger memories now available to store larger programs that are supposed to better support usability. But this has been done without much basis other than individual designer intuition and common sense. Undeniably, common sense, combined with some vigorous trial and error, has already done much to improve systems. Just as undeniably, however, there is much farther to go than we have come. People still complain bitterly about the difficulty of learning to use even text editors and spread sheets, which are among the most thoroughly evolved interactive devices. And there are as yet only a few cognitive tools available that offer totally new ways of accomplishing mental tasks—symbolic math languages are one example—although one would think that the capabilities of computation would open the way for hosts.

Psychologists have become increasingly involved in the design of new computer systems and in research and theory aimed at understanding the human component of the problem. For example, before divestiture, Bell Laboratories alone employed around 200 psychologists in work related to computer system development. Many times that number are employed by other software and hardware companies the world over, so psychology apparently has something to sell. Indeed, talking to applied psychologists and managers in such settings usually elicits tales of success. But tales of dissatisfaction and frustration are also common. It would not be altogether unfair to characterize the situation as follows. Although psychologists have brought to development efforts a dedicated professional interest in user problems, and have successfully acted as intelligent advocates of the interest of users, they have not brought an impressive tool kit of design methods or principles, nor have they effectively brought to bear a relevant body of scientific knowledge or theory.

In addition to helping to build better systems, one would also hope that the interaction of cognitive psychology with design would help to advance the science of mind. Many computer systems are created to interact with, aid, or replace mental processes. Surely the problems encountered in trying to make them do so ought to feed new and interesting psychological research. Computer technology should provide better opportunity for applied research that can contribute to the science of mind than anything we have had in the past. It offers an arena in which potential understanding of human mental powers and limitations can be tested.

The iterative interplay between the invention of new methods to support cognitive activities and the analysis of their successes and failures is a very exciting prospect. In my view, cognitive psychology has suffered from the lack of an applied discipline in which the completeness of its accounts could be measured, or from which a sorting of phenomena into those important for actual human function from those of merely scholastic