

The Human– Computer Interaction Handbook

Fundamentals,
Evolving Technologies,
and Emerging Applications

Third Edition

Edited by Julie A. Jacko, Ph.D.



CRC Press
Taylor & Francis Group



The Human– Computer Interaction Handbook

Fundamentals,
Evolving Technologies,
and Emerging Applications

Third Edition

Human Factors and Ergonomics

Series Editor

Gavriel Salvendy

*Professor Emeritus
School of Industrial Engineering
Purdue University*

*Chair Professor & Head
Dept. of Industrial Engineering
Tsinghua Univ., P.R. China*

PUBLISHED TITLES

- Conceptual Foundations of Human Factors Measurement, *D. Meister*
- Content Preparation Guidelines for the Web and Information Appliances: Cross-Cultural Comparisons
H. Liao, Y. Guo, A. Savoy, and G. Salvendy
- Designing for Accessibility: A Business Guide to Countering Design Exclusion, *S. Keates*
- Handbook of Cognitive Task Design, *E. Hollnagel*
- The Handbook of Data Mining, *N. Ye*
- Handbook of Digital Human Modeling: Research for Applied Ergonomics and Human Factors Engineering
V. G. Duffy
- Handbook of Human Factors and Ergonomics in Health Care and Patient Safety, Second Edition,
P. Carayon
- Handbook of Human Factors in Web Design, Second Edition, *R. Proctor and K. Vu*
- Handbook of Occupational Safety and Health, *D. Koradecka*
- Handbook of Standards and Guidelines in Ergonomics and Human Factors, *W. Karwowski*
- Handbook of Virtual Environments: Design, Implementation, and Applications, *K. Stanney*
- Handbook of Warnings, *M. Wogalter*
- Human-Computer Interaction: Designing for Diverse Users and Domains, *A. Sears and J. A. Jacko*
- Human-Computer Interaction: Design Issues, Solutions, and Applications, *A. Sears and J. A. Jacko*
- Human-Computer Interaction: Development Process, *A. Sears and J. A. Jacko*
- The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies,
and Emerging Applications, Second Edition, *A. Sears and J. A. Jacko*
- Human Factors in System Design, Development, and Testing, *D. Meister and T. Enderwick*
- Introduction to Human Factors and Ergonomics for Engineers, *M. R. Lehto and J. R. Buck*
- Macroergonomics: Theory, Methods and Applications, *H. Hendrick and B. Kleiner*
- Practical Speech User Interface Design, *James R. Lewis*
- Smart Clothing: Technology and Applications, *Gilsoo Cho*
- Theories and Practice in Interaction Design, *S. Bagnara and G. Crampton-Smith*
- The Universal Access Handbook, *C. Stephanidis*
- Usability and Internationalization of Information Technology, *N. Aykin*
- User Interfaces for All: Concepts, Methods, and Tools, *C. Stephanidis*

FORTHCOMING TITLES

- Computer-Aided Anthropometry for Research and Design, *K. M. Robinette*
- Cross-Cultural Design for IT Products and Services, *P. Rau, T. Plocher and Y. Choong*
- Foundations of Human-Computer and Human-Machine Systems, *G. Johannsen*
- The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies,
and Emerging Applications, Third Edition, *J. A. Jacko*
- Handbook of Virtual Environments: Design, Implementation, and Applications, Second Edition
K. S. Hale and K M. Stanney
- Introduction to Human Factors and Ergonomics for Engineers, Second Edition, *M. R. Lehto*
- The Science of Footwear, *R. S. Goonetilleke*
- Skill Training in Multimodal Virtual Environments, *M. Bergamsco, B. Bardy, and D. Gopher*

The Human– Computer Interaction Handbook

Fundamentals,
Evolving Technologies,
and Emerging Applications

Third Edition

Edited by Julie A. Jacko, Ph.D.



CRC Press

Taylor & Francis Group

Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business

Cover art by Torre Schmitt (Torre Schmitt Design, Minneapolis, MN, torreschmitt@gmail.com).

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2012 by Taylor & Francis Group, LLC
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works
Version Date: 20120227

International Standard Book Number-13: 978-1-4398-2944-8 (eBook - PDF)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

This handbook is dedicated to those who have lent a hand and lit the way.

And I said to the man who stood at the gate of the year:

“Give me a light, that I may tread safely into the unknown!”

And he replied:

“Go out into the darkness and put your hand into the Hand of God.

That shall be to you better than light and safer than a known way.”

*King George VI in his New Year’s message to his embattled people
at the beginning of the Second World War*

Contents

Series Foreword	xiii
<i>Gavriel Salvendy</i>	
Foreword	xv
<i>Ben Shneiderman</i>	
Preface.....	xvii
<i>Julie A. Jacko</i>	
Editor	xix
Contributors	xxi

Introduction A Moving Target: The Evolution of Human–Computer Interaction.....	xxvii
<i>Jonathan Grudin</i>	

PART I Humans in HCI

Chapter 1 Perceptual-Motor Interaction: Some Implications for Human–Computer Interaction	3
<i>Timothy N. Welsh, Sanjay Chandrasekharan, Matthew Ray, Heather Neyedli, Romeo Chua, and Daniel J. Weeks</i>	
Chapter 2 Human Information Processing: An Overview for Human–Computer Interaction.....	21
<i>Robert W. Proctor and Kim-Phuong L. Vu</i>	
Chapter 3 Mental Models in Human–Computer Interaction.....	41
<i>Stephen J. Payne</i>	
Chapter 4 Task Loading and Stress in Human–Computer Interaction: Theoretical Frameworks and Mitigation Strategies	55
<i>James L. Szalma, Gabriella M. Hancock, and Peter A. Hancock</i>	
Chapter 5 Choices and Decisions of Computer Users	77
<i>Anthony Jameson</i>	

PART II Computers in HCI

Chapter 6 Input Technologies and Techniques	95
<i>Ken Hinckley and Daniel Wigdor</i>	
Chapter 7 Sensor- and Recognition-Based Input for Interaction	133
<i>Andrew D. Wilson</i>	

Chapter 8	Visual Displays.....	157
	<i>Christopher M. Schlick, Carsten Winkelholz, Martina Ziefle, and Alexander Mertens</i>	
Chapter 9	Haptic Interface.....	193
	<i>Hiroo Iwata</i>	
Chapter 10	Nonspeech Auditory and Crossmodal Output	211
	<i>Eve Hoggan and Stephen Brewster</i>	
Chapter 11	Network-Based Interaction.....	237
	<i>Alan Dix</i>	
Chapter 12	Wearable Computers	273
	<i>Daniel Siewiorek, Asim Smailagic, and Thad Starner</i>	
Chapter 13	Design of Fixed, Portable, and Mobile Information Devices	297
	<i>Michael J. Smith and Pascale Carayon</i>	

PART III *Designing Human–Computer Interactions*

Chapter 14	Visual Design Principles for Usable Interfaces: Everything Is Designed: Why We Should Think before Doing.....	315
	<i>Suzanne Watzman and Margaret Re</i>	
Chapter 15	Globalization, Localization, and Cross-Cultural User-Interface Design.....	341
	<i>Aaron Marcus and Emilie W. Gould</i>	
Chapter 16	Speech and Language Interfaces, Applications, and Technologies	367
	<i>Clare-Marie Karat, Jennifer Lai, Osamuyimen Stewart, and Nicole Yankelovich</i>	
Chapter 17	Multimedia User Interface Design	387
	<i>Alistair Sutcliffe</i>	
Chapter 18	Multimodal Interfaces	405
	<i>Sharon Oviatt</i>	
Chapter 19	Systems That Adapt to Their Users.....	431
	<i>Anthony Jameson and Krzysztof Z. Gajos</i>	
Chapter 20	Mobile Interaction Design in the Age of Experience Ecosystems.....	457
	<i>Marco Susani</i>	
Chapter 21	Tangible User Interfaces.....	465
	<i>Hiroshi Ishii and Brygg Ullmer</i>	

Chapter 22	Achieving Psychological Simplicity: Measures and Methods to Reduce Cognitive Complexity.....	491
	<i>John C. Thomas and John T. Richards</i>	
Chapter 23	Information Visualization	515
	<i>Stuart Card</i>	
Chapter 24	Collaboration Technologies.....	549
	<i>Gary M. Olson and Judith S. Olson</i>	
Chapter 25	Human–Computer Interaction and the Web	565
	<i>Helen Ashman, Declan Dagger, Tim Brailsford, James Goulding, Declan O’Sullivan, Jan-Felix Schmakeit, and Vincent Wade</i>	
Chapter 26	Human-Centered Design of Decision-Support Systems	589
	<i>Philip J. Smith, Roger Beatty, Caroline C. Hayes, Adam Larson, Norman D. Geddes, and Michael C. Dorneich</i>	
Chapter 27	Online Communities	623
	<i>Panayiotis Zaphiris, Chee Siang Ang, and Andrew Laghos</i>	
Chapter 28	Virtual Environments.....	643
	<i>Kay M. Stanney and Joseph V. Cohn</i>	
Chapter 29	Privacy, Security, and Trust: Human–Computer Interaction Challenges and Opportunities at Their Intersection.....	669
	<i>John Karat, Clare-Marie Karat, and Carolyn Brodie</i>	
 PART IV <i>Application-/Domain-Specific Design</i>		
Chapter 30	Human–Computer Interaction in Health Care.....	701
	<i>François Sainfort, Julie A. Jacko, Molly A. McClellan, and Paula J. Edwards</i>	
Chapter 31	Why We Play: Affect and the Fun of Games—Designing Emotions for Games, Entertainment Interfaces, and Interactive Products	725
	<i>Nicole Lazzaro</i>	
Chapter 32	Motor Vehicle–Driver Interfaces	749
	<i>Paul A. Green</i>	
Chapter 33	Human–Computer Interaction in Aerospace	771
	<i>Steven J. Landry</i>	

Chapter 34	User-Centered Design in Games	795
	<i>Randy J. Pagulayan, Kevin Keeker, Thomas Fuller, Dennis Wixon, Ramon L. Romero, and Daniel V. Gunn</i>	

PART V *Designing for Diversity*

Chapter 35	Older Adults and Information Technology: Opportunities and Challenges	825
	<i>Sara J. Czaja and Chin Chin Lee</i>	
Chapter 36	Human–Computer Interaction for Kids	841
	<i>Amy Bruckman, Alisa Bandlow, Jill Dimond, and Andrea Forte</i>	
Chapter 37	Information Technology for Communication and Cognitive Support	863
	<i>Alan F. Newell, Alex Carmichael, Peter Gregor, Norman Alm, Annalu Waller, Vicki L. Hanson, Graham Pullin, and Jesse Hoey</i>	
Chapter 38	Perceptual Impairments: New Advancements Promoting Technological Access	893
	<i>Julie A. Jacko, V. Kathlene Leonard, Molly A. McClellan, and Ingrid U. Scott</i>	
Chapter 39	Universal Accessibility and Low-Literacy Populations: Implications for Human–Computer Interaction Design and Research Methods	913
	<i>William M. Gribbons</i>	
Chapter 40	Computing Technologies for Deaf and Hard of Hearing Users	933
	<i>Vicki L. Hanson</i>	

PART VI *The Development Process*

SECTION A *Requirements Specification*

Chapter 41	User Experience Requirements Analysis within the Usability Engineering Lifecycle	945
	<i>Deborah J. Mayhew and Todd J. Follansbee</i>	
Chapter 42	Task Analysis	955
	<i>Catherine Courage, Jhilmil Jain, Janice (Ginny) Redish, and Dennis Wixon</i>	
Chapter 43	Contextual Design	983
	<i>Karen Holtzblatt</i>	
Chapter 44	Grounded Theory Method in Human–Computer Interaction and Computer-Supported Cooperative Work	1003
	<i>Michael J. Muller and Sandra Kogan</i>	
Chapter 45	An Ethnographic Approach to Design	1025
	<i>Jeanette Blomberg and Mark Burrell</i>	

SECTION B *Design and Development*

Chapter 46	Putting Personas to Work: Employing User Personas to Focus Product Planning, Design, and Development	1055
	<i>John Pruitt and Tamara Adlin</i>	
Chapter 47	Prototyping Tools and Techniques	1081
	<i>Michel Beaudouin-Lafon and Wendy E. Mackay</i>	
Chapter 48	Scenario-Based Design	1105
	<i>Mary Beth Rosson and John M. Carroll</i>	
Chapter 49	Participatory Design: The Third Space in Human–Computer Interaction	1125
	<i>Michael J. Muller and Allison Druin</i>	
Chapter 50	Unified User Interface Development: A Software Refactoring Perspective	1155
	<i>Anthony Savidis and Constantine Stephanidis</i>	
Chapter 51	Usability + Persuasiveness + Graphic Design = eCommerce User Experience	1181
	<i>Deborah J. Mayhew</i>	
Chapter 52	Human–Computer Interaction and Software Engineering for User Interface Plasticity	1195
	<i>Joëlle Coutaz and Gaëlle Calvary</i>	

SECTION C *Testing, Evaluation, and Technology Transfer*

Chapter 53	Usability Testing.....	1221
	<i>Joseph S. Dumas and Jean E. Fox</i>	
Chapter 54	Usability for Engaged Users: The Naturalistic Approach to Evaluation.....	1243
	<i>David Siegel</i>	
Chapter 55	Survey Design and Implementation in HCI	1259
	<i>A. Ant Ozok</i>	
Chapter 56	Inspection-Based Evaluations	1279
	<i>Gilbert Cockton, Alan Woolrych, Kasper Hornbæk, and Erik Frøkjær</i>	
Chapter 57	Model-Based Evaluation	1299
	<i>David Kieras</i>	
Chapter 58	Spreadsheet Tool for Simple Cost-Benefit Analyses of User Experience Engineering	1319
	<i>Deborah J. Mayhew</i>	

Chapter 59 Technology Transfer	1331
<i>Kevin M. Schofield</i>	

PART VII Emerging Phenomena in HCI

Chapter 60 Augmenting Cognition in HCI: Twenty-First Century Adaptive System Science and Technology	1343
<i>Kelly S. Hale, Kay M. Stanney, and Dylan D. Schmorrow</i>	

Chapter 61 Social Networks and Social Media	1359
<i>Molly A. McClellan, Julie A. Jacko, François Sainfort, and Layne M. Johnson</i>	

Chapter 62 Human–Computer Interaction for Development: Changing Human–Computer Interaction to Change the World	1375
<i>Susan M. Dray, Ann Light, Andrew M. Dearden, Vanessa Evers, Melissa Densmore, Divya Ramachandran, Matthew Kam, Gary Marsden, Nithya Sambasivan, Thomas Smyth, Darelle van Greunen, and Niall Winters</i>	

Series Foreword

The third edition of this classic handbook is published at an opportune time when interactive technologies are a dominating presence in work, leisure, and social settings and when ambient intelligence is gaining accelerated momentum. The field of human–computer interaction (HCI) has matured to such an extent that even the words comprising the term have taken on new, expanded, and reinterpreted meanings. That is, the field has advanced significantly from its origins. Researchers in HCI are called upon now more than ever to develop new knowledge, which often resides at the intersection of multiple disciplines and spans various and innovative platforms of applications. Information technology is more ubiquitous today than ever, successfully interacting with the technologies that ensure it is more enjoyable and more productively accessible and usable by all segments of society across all five continents.

This handbook is the premier resource for the theoretical and operational foundations of HCI, providing readers access to the latest scientific breakthroughs coupled with the state of the art in the field. The book provides detailed descriptions of approaches and methodologies that are frequently illustrated with case studies and examples on how to conceptualize,

design, and evaluate interactive systems with human beings at the center of the endeavor. As such, this handbook will be invaluable to researchers, practitioners, educators, and students working in, or at the intersection of, computer science, information technology, information science, informatics, engineering, psychology, design, and human factors and ergonomics.

This book is part of the Human Factors and Ergonomics series, published by the Taylor & Francis Group. The 145 authors of this handbook include 92 from academia, 49 from industry, and 4 from government agencies. These individuals are among the very best and most respected in their fields across the globe. The more than 80 tables, 400 figures, and nearly 7000 references in this book provide the single most comprehensive depiction of this field that exists in a single volume.

The handbook authors come from 14 countries: Australia, Canada, Cyprus, Denmark, Finland, France, Germany, Greece, Ireland, Japan, South Africa, the Netherlands, United Kingdom, and the United States.

Gavriel Salvendy, Series Editor

Purdue University/Tsinghua University, China

Foreword

The Expanding Impact of Human–Computer Interaction

The remarkable growth of human–computer interaction (HCI) over the past 30 years has transformed this nascent interdisciplinary field into an intellectually rich and high impact worldwide phenomenon. We have grown from a small rebellious group of researchers who struggled to gain recognition as we broke disciplinary boundaries to a broad influential community with potent impact on the daily lives of every human. There are dozens of relevant journals, plus conferences and workshops worldwide.

The aspirations of early HCI researchers and practitioners were to make better menus, design graphical user interfaces based on direct manipulation, improve input devices, design effective control panels, and present information in comprehensible formats. HCI software developers contributed innovative tools that enabled programmers and nonprogrammers to create interfaces for widely varying applications and diverse users. HCI professionals developed design principles, guidelines, and sometimes standards dealing with consistency, informative feedback, error prevention, shortcuts for experts, and user control. Success was measured by individual performance metrics such as learning time, speed, error rates, and retention for specific tasks, whereas user satisfaction was assessed by detailed questionnaires filled with numbered scales.

In the early days, HCI researchers and professionals fought to gain recognition and often still have to justify HCI's value with academic colleagues or corporate managers. However, the larger world embraced our contributions and now has high expectations of what we can deliver. Few fields can claim such rapid expansion and broad impact as those who design the desktop, web, mobile, and cellphone interfaces that have spread around the world into the hands of at least 5 billion users. HCI designs now influence commercial success, reform education, change family life, affect the political stability of nations, are embedded in military systems and play a significant role in shaping a peaceful or conflict-ridden world.

The Handbook of Human–Computer Interaction: Third Edition details the progress of this extraordinary discipline, inviting newcomers to learn about it and helping experienced professionals to understand the rapid and continuing changes. The carefully written chapters and extensive references will be useful to readers who want to scan the territory or dig deep into specific topics. This handbook's prominent authors thoughtfully survey the key topics, enabling students, researchers, and professionals to appreciate HCI's impact.

As HCI progresses, there is a greater acceptance in the academic environment, where HCI is now part of most computer science, iSchool, business, engineering, and other departments and has advocates in medicine, social sciences, journalism, humanities, etc. Although the term *human–computer interaction* has achieved widespread recognition, many insiders feel that it is no longer an accurate description. They complain that it suggests one human interacting with one computer to complete narrow tasks. Instead, these critics believe that the discipline should reflect user-oriented technologies that are ubiquitous, pervasive, social, embedded, tangible, invisible, multimodal, immersive, augmented, or ambient. Some want to break free from the focus on computer use and emphasize user experiences, interaction design, emotional impact, aesthetics, social engagement, empathic interactions, trust building, and human responsibility.

New terms have been proposed such as *human-centered computing*, *social computing*, *human–information interaction*, *human–social interaction*, *human-centered informatics*, or just *human interaction*. Novel, but already thriving applications areas include computational biology, computational social science, e-commerce (and m-commerce), digital humanities, information visualization, open government, sustainability, biodiversity, and citizen science. Although these broader visions are important, many researchers are still working on innovative display designs, input devices, multimedia output, programming toolkits, and predictive models of user performance.

New names and applications are a good sign of success, but finding the balance between sticking with an established term and welcoming innovative directions is difficult. Maybe an old aphorism helps: “make new friends and keep the old, one is silver and the other gold.” Can we retain the brand name recognition of HCI but embrace new directions by discussing *micro-HCI* and *macro-HCI*?

Micro-HCI researchers and developers would design and build innovative interfaces and deliver validated guidelines for use across the range of desktop, web, mobile, and ubiquitous devices. The challenges for micro-HCI are to deal with rapidly changing technologies, while accommodating the wide range of users: novice/expert, young/old, literate/illiterate, abled/disabled, and their cultural plus linguistic diversity. These distinctions are tied to skills, but there are further diversities in gender, personality, ethnicity, skills, and motivation that are now necessary to address in interface designs. Micro-HCI researchers can take comfort in dealing

with well-stated requirements, clear benchmark tasks, and effective predictive models.

Macro-HCI researchers and developers would explore new design territories such as affective experience, aesthetics, motivation, social participation, trust, empathy, responsibility, and privacy. The challenges for macro-HCI are to deal with new opportunities across the range of human experience: commerce, law, health/wellness, education, creative arts, community relationships, politics, policy negotiation, conflict resolution, international development, and peace studies. Macro-HCI researchers have to face the challenge of more open tasks, unanticipated user goals, and even conflicts among users in large communities.

Although micro-HCI and macro-HCI have healthy overlaps, as do micro-economics and macro-economics, they attract different types of researchers, practitioners, and activists, thereby further broadening the scope and impact. As commercial, social, legal, and ethical considerations play an increasing role, educational curricula and professional practices need to be updated regularly and midcareer continuing education for HCI professionals will keep them current.

An important goal will be to develop new metrics and evaluation methods for micro-HCI and macro-HCI. Moore's Law has been useful in charting the growth of computing, enabling everyone to admire and benefit from the increase in gigahertz, terabytes, and petaflops. These are still useful, but we need newer metrics to understand the impact of HCI designs that have enabled the spread of billions of mobile devices and the emergence of YouTube, Facebook, twitter, Wikipedia, and so on. Understanding this transformation would be facilitated by measures of giga-hellos, tera-contribs, and peta-thankyous and by newer metrics of trust, empathy, responsibility, privacy, and so on.

Traditional evaluation approaches of controlled experiments and usability testing are being continuously refined to fit the needs of micro-HCI, whereas the newer methods of qualitative, ethnographic, and case study methods are being explored to match the needs of macro-HCI. Both groups will benefit from the remarkable increased opportunities to log usage on a massive scale through the increasingly connected communications, data, and sensor networks. Traditional surveys of a small sample of users who offer biased perceptions or reports of attitudes are giving way to actual measurement

of usage that reveals the learnability, efficacy, utility, and satisfaction of users. Even more exciting is the potential to capture the manifestations of trust, empathy, responsibility, privacy, security, and motivation. Researchers are also beginning to measure brand loyalty, parental engagement, political leaning, potential for violence, community commitment, and much more. The dangers of inappropriate intrusion, misguided applications, scamming/spamming, deception, and bullying are now part of macro-HCI. Even greater concerns come from criminals, terrorists, and oppressive governments who can use these technologies in ways that threaten individuals, intimidate communities, or destroy the environment.

The power of widely used social technologies that stem from HCI's success means that we will face ethical challenges similar to what the nuclear physicists dealt with during the 1940s and beyond. We cannot and should not avoid these responsibilities. Rather, we should embrace them and show leadership in shaping technology to produce positive outcomes. This is never easy, but every worthy project that improves the health, environment, or education of children or builds capacity for constructive communities should be recognized, disseminated, scaled up, and continuously improved. Even more ambitious should be our efforts to promote open government, independent oversight, deliberative systems, and citizen participation. The research agenda for HCI should include the UN Millennium Development Goals such as eradicating extreme hunger and poverty, ensuring universal childhood education, promoting maternal health, and ensuring environmental sustainability. If HCI professionals also courageously address conflict resolution, international development, and peace studies, we can inspire others and help build a better world.

We should be proud of what HCI has accomplished, but there is much work to be done. Let's get on with it!

ACKNOWLEDGMENTS

Thanks to Ron Baecker, Jack Carroll, Susan Dray, Gerhard Fischer, Rob Jacob, Clare-Marie Karat, Clayton Lewis, Brad Myers, and Jenny Preece for comments on earlier drafts.

Ben Shneiderman
University of Maryland

Preface

This third edition of the HCI handbook represents the single largest, most complete compilation of HCI theories, principles, advances, case studies, and more that exist within a single volume. The construction of the handbook has been a massive community effort of which it was a tremendous privilege for this author and editor to be a part. The 145 authors of the 62 chapters within this book are people who have not only dedicated themselves to laying the foundation for this field but also dared to address the grand challenges that have been posed along the way, thus advancing the field of HCI by leaps and bounds. The HCI community from which these authors hail is remarkably diverse and collaborative. You will see the artifacts of this ethos throughout the book.

The handbook opens with an insightful and thought-provoking introduction written by Jonathan Grudin, which sets the tone for the entire book. Within the introduction you will find a unique and compelling depiction of the evolution of HCI. The handbook closes with a look at the evolving nature of HCI to change the world. The closing chapter is written by the largest collection of authors in the book, led by Susan Dray. The global focus of this chapter is personified by the authors' origins, which literally span the globe. The chapters in between are organized very much like those in the second edition; however, the content of the chapters has been dramatically updated to reflect the state of the art and current state of the science in HCI. There have been numerous notable additions to the third edition, which reflect the ever-growing nature of this field, including, for example, chapters on social networks and social media, grounded theory, choices and decisions of users, and the naturalistic approach to evaluation.

I offer my heartfelt thanks to Ben Shneiderman, who kindly agreed to contribute his revolutionary perspective in the Foreword to the third edition. He not only chronicles the impact of HCI but also presents a challenge to each and every one of us to embrace the responsibility of shaping technology

to produce positive outcomes. With this challenge he is asking us to be the best citizen scholars we can be. This is classic Ben Shneiderman and just one of the many reasons why I respect and admire him. This handbook would simply not have been possible without the guiding influence of my longtime mentor and good friend, Gavriel Salvendy. Gavriel sets the standard for successfully coalescing people and communities around shared goals and mutual aspirations. He has been an unwavering source of inspiration, support, advice, opportunity, and kindness for me. This book is part of a larger book series of which Gavriel is the series editor. His Series Foreword to the third edition enables us to see this book in the context of the larger whole. Both these luminaries, Ben and Gavriel, have transformed the field of HCI in their own signature ways, and I salute both of them.

A very special individual worked hand in hand with me in constructing the third edition. Molly McClellan, PhD, is a research associate with SimPORTAL at the University of Minnesota, performing postdoctoral research in the area of perioperative simulation. Completing a book of this scale and scope requires incredible persistence and perseverance. Molly demonstrates both these attributes and so much more. She is a creative problem solver with an uncanny ability to organize vast quantities of information from disparate and geographically distributed sources. She is smart, generous, and exceedingly committed to excellence. I have admired her as a scholar and as a human being. It is a privilege to serve as her major professor and mentor.

Last but not the least, I wish to recognize the support offered me by my husband François and our son Nico. They are both, quite simply, my *raison de vivre*.

Julie A. Jacko
University of Minnesota

Editor

Julie A. Jacko, PhD, is a professor of Public Health at the University of Minnesota and a faculty fellow in the Academic Health Center's Institute for Health Informatics. She is the principal investigator and director of the University Partnership for Health Informatics (UP-HI). This \$5.1-million grant from the Office of the National Coordinator for Health Information Technology is one of the nine awarded nationally and represents the first public-private partnership funded in the upper Midwest United States to infuse our nation's workforce with individuals who have been trained to perform one of six mission-critical health information technology roles.

Dr. Jacko has expertise in the design, implementation, and evaluation of interactive computing systems in complex domains such as population health and health care delivery, with the purpose of enhancing human performance and satisfaction. This is accomplished through research that is focused on the cognitive processes underlying the interaction of people with complex systems with the ultimate goal of combining robust empirical results with the development of engineering models of human performance that can aid in the design of real-world systems. Dr. Jacko has an exemplary research track record spanning nearly 20 years during which over 160 scientific publications have been generated in these research areas. She has generated nearly \$25 million by way of research funding in the last 10 years and is one of the only 20 recipients of a National Science Foundation Presidential Early Career Award for Scientists and Engineers, the highest honor awarded to young investigators by the U.S. government. Dr. Jacko served as co-author on the #1-rated published article for 2005 in the *International Journal of Medical Informatics*. She has an extensive track record of professional leadership excellence, including the following awards and honorable positions.

- Received commendation from the Office of the National Coordinator for Health Information Technology (ONC) for her innovation and leadership in the University Partnership for Health Informatics (2011).

- Ranked one of the Top Ten Influential Informatics Professors by HealthTechTopia on September 14, 2010. (<http://mastersinhealthinformatics.com/2010/top-10-most-influential-informatics-professors/>).
- Published an invited comment paper in 2011, Issue 470, in *Nature* titled, "Narrow the Gap in Health Literacy."
- Appointed by the State of Minnesota Commissioner of Health to serve on the Minnesota e-health Advisory Committee, representing academics and clinical research for the State of Minnesota.
- Elected to the Office of the President, the Association for Computing Machinery's Special Interest Group on Computer-Human Interaction (ACM SIGCHI) (2006-09).
- Elected to the Association for Computing Machinery Special Interest Group Governing Board Executive Committee, elected Member at Large (2007-2009).
- Elected to the Office of the Vice President for Membership and Communications—the Association for Computing Machinery's Special Interest Group on Computer-Human Interaction (ACM SIGCHI) (2003-06).
- Editor-in-chief of the *International Journal of Human-Computer Interaction*, published by Taylor & Francis.
- Co-editor of the 1st and 2nd Editions of the *Human-Computer Interaction Handbook*, the premier compendium of research and practice in the field of human-computer interaction.

In addition, during the last 15 years, Dr. Jacko has chaired or co-chaired numerous technical conferences and technical conference programs in the fields of human factors and human-computer interaction. She received her BS, MS, and PhD in industrial and systems engineering from Purdue University in West Lafayette, Indiana, where she held the NEC Graduate Fellowship.

Contributors

Tamara Adlin

Adlin, Inc.
Seattle, Washington

Norman Alm

School of Computing
University of Dundee
Dundee, Scotland

Chee Siang Ang

School of Engineering and Digital Arts
University of Kent
Canterbury, United Kingdom

Helen Ashman

School of Computer and Information
Science
University of South Australia
Adelaide, Australia

Alisa Bandlow

Sandia National Laboratories
Albuquerque, New Mexico

Roger Beatty

Operational Control Services
Dallas/Fort Worth, Texas

Michel Beaudouin-Lafon

Department of Computer Science
Université Paris-Sud
Paris, France

Jeanette Blomberg

IBM Almaden Research Center
San Jose, California

Tim Brailsford

School of Computer Science and
Information Technology
University of Nottingham
Nottingham, United Kingdom

Stephen Brewster

Department of Computing Science
University of Glasgow
Glasgow, Scotland

Carolyn Brodie

Brodie Consulting Group
Indianola, Iowa

Amy Bruckman

College of Computing
Georgia Institute of Technology
Atlanta, Georgia

Mark Burrell

Endeca Technologies
Cambridge, Massachusetts

Gaëlle Calvary

Grenoble Informatics Laboratory
Institut Polytechnique de Grenoble
Grenoble, France

Pascale Carayon

Department of Industrial and Systems
Engineering
Center for Quality and Productivity
Improvement (CQPI)
University of Wisconsin
Madison, Wisconsin

Stuart Card

Xerox PARC
Palo Alto, California

Alex Carmichael

School of Computing
University of Dundee
Dundee, Scotland

John M. Carroll

College of Information Sciences and
Technology
Pennsylvania State University
University Park, Pennsylvania

Sanjay Chandrasekharan

School of Interactive Computing
Georgia Institute of Technology
Atlanta, Georgia

Romeo Chua

School of Human Kinetics
University of British Columbia
Vancouver, British Columbia, Canada

Gilbert Cockton

School of Design
Northumbria University
New Castle, United Kingdom

Joseph V. Cohn

Office of Naval Research
Arlington, Virginia

Catherine Courage

Citrix Systems
Fort Lauderdale, Florida

Joëlle Coutaz

Grenoble Informatics Laboratory
Université Joseph Fourier
Grenoble, France

Sara J. Czaja

Department of Psychiatry and
Behavioral Sciences
Center on Aging
University of Miami Miller School of
Medicine
Miami, Florida

Declan Dagger

Department of Computer Science
Trinity College
Dublin, Ireland

Andrew M. Dearden

Communication and Computing
Research Centre
Sheffield Hallam University
South Yorkshire, United Kingdom

Melissa Densmore

School of Information
University of California
Berkeley, California

Jill Dimond

College of Computing
Georgia Institute of Technology
Atlanta, Georgia

Alan Dix

Talis
Birmingham, United Kingdom

Michael C. Dorneich

Human-Centered Systems
Honeywell Laboratories, Golden Valley
Minneapolis, Minnesota

Susan M. Dray

Dray and Associates, Inc.
Minneapolis, Minnesota

Allison Druin

Human–Computer Interaction Lab
College of Information Studies
University of Maryland
College Park, Maryland

Joseph S. Dumas

Dumas Consulting
Yarmouth Port, Massachusetts

Paula J. Edwards

HIMFormatics, LLC
Atlanta, Georgia

Vanessa Evers

Department of Electrical Engineering
and Computer Science
University of Twente
Enschede, the Netherlands

Todd J. Follansbee

Web Marketing Resource
West Tisbury, Massachusetts

Andrea Forte

College of Information Science and
Technology
Drexel University
Philadelphia, Pennsylvania

Jean E. Fox

Bureau of Labor Statistics
Washington, D.C.

Erik Frøkjær

Department of Computer Science
University of Copenhagen
Copenhagen, Denmark

Thomas Fuller

Microsoft Studios
Redmond, Washington

Krzysztof Z. Gajos

Harvard School of Engineering and
Applied Sciences
Harvard University
Cambridge, Massachusetts

Norman D. Geddes

Applied Systems Intelligence, Inc.
Alpharetta, Georgia

Emilie W. Gould

Department of Communication
State University of New York
at Albany
Albany, New York

James Goulding

School of Computer Science and
Information Technology
University of Nottingham
Nottingham, United Kingdom

Paul A. Green

Driver Interface Group
Transportation Research Institute
University of Michigan
Ann Arbor, Michigan

Peter Gregor

School of Computing
University of Dundee
Dundee, Scotland

William M. Gribbons

Department of Human Factors
Bentley University
Waltham, Massachusetts

Jonathan Grudin

Microsoft Research
Redmond, Washington

Daniel V. Gunn

Microsoft Studios
Redmond, Washington

Kelly S. Hale

Design Interactive, Inc.
Oviedo, Florida

Gabriella M. Hancock

Department of Applied Physiology and
Kinesiology
University of Florida
Gainesville, Florida

Peter A. Hancock

Department of Psychology
and
Institute for Simulation and Training
University of Central Florida
Orlando, Florida

Vicki L. Hanson

School of Computing
University of Dundee
Dundee, Scotland

Caroline C. Hayes

Department of Mechanical
Engineering
University of Minnesota
Minneapolis/St. Paul, Minnesota

Ken Hinckley

Microsoft Research
Redmond, Washington

Jesse Hoey

School of Computer Science
University of Waterloo
Waterloo, Ontario, Canada

Eve Hoggan

Department of Computer Science
University of Helsinki
Helsinki, Finland

Karen Holtzblatt

InContext Design
Concord, Massachusetts

Kasper Hornbæk

Department of Computer Science
University of Copenhagen
Copenhagen, Denmark

Hiroshi Ishii

MIT Media Lab
Massachusetts Institute of Technology
Cambridge, Massachusetts

Hiroo Iwata

Graduate School of Systems and
Information Engineering
University of Tsukuba
Tsukuba, Japan

Julie A. Jacko

School of Public Health Division of
Environmental and Health Sciences
University of Minnesota
Minneapolis, Minnesota

Jhilmil Jain

Google
Mountain View, California

Anthony Jameson

DFKI, German Research Center for
Artificial Intelligence
Saarbrücken, Germany

Layne M. Johnson

Health Sciences Library and Institute
for Health Informatics
University of Minnesota
Minneapolis, Minnesota

Matthew Kam

Human–Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania

Clare-Marie Karat

Karat Consulting Group
Aptos, California

John Karat

Karat Consulting Group
Aptos, California

Kevin Keeker

Zynga, Inc.
San Francisco, California

David Kieras

Electrical Engineering and Computer
Science Department
University of Michigan
Ann Arbor, Michigan

Sandra Kogan

IBM Corporation
Cambridge, Massachusetts

Andrew Laghos

Department of Multimedia and
Graphic Arts
Cyprus University of Technology
Lemesos, Cyprus

Jennifer Lai

IBM T. J. Watson Research Center
Yorktown Heights, New York

Steven J. Landry

School of Industrial Engineering
Purdue University
West Lafayette, Indiana

Adam Larson

United States Air Force

Nicole Lazzaro

XEODesign, Inc.
Oakland, California

Chin Chin Lee

Center on Aging
University of Miami Miller School of
Medicine
Miami, Florida

V. Kathlene Leonard

Cloverture, LLC
Smyrna, Georgia

Ann Light

Communication and Computing
Research Centre
Sheffield Hallam University
Sheffield, United Kingdom

Wendy E. Mackay

INRIA Saclay
Orsay, France

Aaron Marcus

Aaron Marcus and Associates, Inc.
Berkeley, California

Gary Marsden

Department of Computer Science
University of Cape Town
Cape Town, South Africa

Deborah J. Mayhew

Deborah J. Mayhew & Associates
West Tisbury, Massachusetts

Molly A. McClellan

SimPORTAL Medical School
University of Minnesota
Minneapolis, Minnesota

Alexander Mertens

Institute of Industrial Engineering and
Ergonomics
RWTH Aachen University
Aachen, Germany

Michael J. Muller

IBM Research
Cambridge, Massachusetts

Alan F. Newell

School of Computing
University of Dundee
Dundee, Scotland

Heather Neyedli

Faculty of Kinesiology and Physical
Education
University of Toronto
Toronto, Ontario, Canada

Gary M. Olson

School of Information and Computer
Sciences
University of California
Irvine, California

Judith S. Olson

School of Information and Computer
Sciences
University of California
Irvine, California

Declan O’Sullivan

Department of Computer Science
Trinity College
Dublin, Ireland

Sharon Oviatt

Incaa Designs
Seattle, Washington

A. Ant Ozok

Department of Information
Systems
University of Maryland, Baltimore
County
Baltimore, Maryland

Randy J. Pagulayan

Microsoft Studios
Redmond, Washington

Stephen J. Payne

University of Manchester
Manchester, United Kingdom

Robert W. Proctor

Department of Psychological
Sciences
Purdue University
West Lafayette, Indiana

John Pruitt

Microsoft Corporation
Redmond, Washington

Graham Pullin

School of Computing
University of Dundee
Dundee, Scotland

Divya Ramachandran

Computer Science Division
and
Berkeley Institute of Design
University of California
Berkeley, California

Matthew Ray

Faculty of Kinesiology and Physical
Education
University of Toronto
Toronto, Ontario, Canada

Margaret Re

Visual Arts Department
University of Maryland, Baltimore
County
Baltimore, Maryland

Janice (Ginny) Redish

Redish & Associates, Inc.
Bethesda, Maryland

John T. Richards

IBM T. J. Watson Research
Center
Yorktown Heights, New York

Ramon L. Romero

Interactive Entertainment
Business
Microsoft Corporation
Redmond, Washington

Mary Beth Rosson

College of Information Sciences and
Technology
Pennsylvania State University
University Park, Pennsylvania

François Sainfort

School of Public Health
University of Minnesota
Minneapolis, Minnesota

Nithya Sambasivan

Department of Informatics
University of California
Irvine, California

Anthony Savidis

ICS-FORTH and Department of
Computer Science
University of Crete
Crete, Greece

Christopher M. Schlick

Institute of Industrial Engineering
and Ergonomics
RWTH Aachen University
Aachen, Germany
and

Fraunhofer-Institute for
Communication, Information
Processing, and Ergonomics
Wachtberg, Germany

Jan-Felix Schmakeit

School of Computer and Information
Science
University of South Australia
Adelaide, Australia

Dylan D. Schmorow

U.S. Navy
Arlington, Virginia

Kevin M. Schofield

Microsoft Research
Redmond, Washington

Ingrid U. Scott

Pennsylvania State College of Medicine
University Park, Pennsylvania

David Siegel

Dray and Associates, Inc.
Minneapolis, Minnesota

Daniel Siewiorek

Human–Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania

Asim Smailagic

Institute for Complex Engineered
Systems
Carnegie Mellon University
Pittsburgh, Pennsylvania

Michael J. Smith

Department of Industrial and Systems
Engineering
Center for Quality and Productivity
Improvement
University of Wisconsin
Madison, Wisconsin

Philip J. Smith

Cognitive Systems Engineering
Laboratory
The Ohio State University
Columbus, Ohio

Thomas Smyth

School of Interactive Computing
Georgia Institute of Technology
Atlanta, Georgia

Kay M. Stanney

Design Interactive, Inc.
Oviedo, Florida

Thad Starner

College of Computing
Georgia Institute of Technology
Atlanta, Georgia

Constantine Stephanidis

ICS-FORTH and Department of
Computer Science
University of Crete
Crete, Greece

Osamuyimen Stewart

IBM T. J. Watson Research Center
Yorktown Heights, New York

Marco Susani

Koz Susani Design
Chicago, Illinois

Alistair Sutcliffe

University of Manchester
Manchester, United Kingdom

James L. Szalma

Department of Psychology
University of Central Florida
Orlando, Florida

John C. Thomas

IBM T. J. Watson Research Center
Yorktown Heights, New York

Brygg Ullmer

Department of Computer Science
and
Center for Computation and
Technology
Louisiana State University
Baton Rouge, Louisiana

Darelle van Greunen

Department of Information Technology
School of Information Technology
Nelson Mandela University
Port Elizabeth, South Africa

Kim-Phuong L. Vu

Department of Psychology
California State University
Long Beach, California

Vincent Wade

School of Computer Science and
Statistics
Trinity College
Dublin, Ireland

Annalu Waller

School of Computing
University of Dundee
Dundee, Scotland

Suzanne Watzman

Watzman Information Design
Somerville, Massachusetts

Daniel J. Weeks

Department of Psychology
University of Lethbridge
Lethbridge, Alberta, Canada

Timothy N. Welsh

Faculty of Kinesiology and Physical
Education
University of Toronto
Toronto, Ontario, Canada

Daniel Wigdor

Department of Computer
Science
University of Toronto
Toronto, Ontario, Canada

Andrew D. Wilson

Microsoft Research
Redmond, Washington

Carsten Winkelholz

Fraunhofer-Institute for
Communication, Information
Processing and Ergonomics
Wachtberg, Germany

Niall Winters

London Knowledge Lab
Institute of Education
University of London
London, United Kingdom

Dennis Wixon

Startup Business Group
Microsoft Corporation
Redmond, Washington

Alan Woolrych

Department of Computing,
Engineering, and Technology
University of Sunderland
United Kingdom

Nicole Yankelovich

Open Wonderland Foundation
Weston, Massachusetts

Panayiotis Zaphiris

Department of Multimedia and
Graphic Arts
Cyprus University of
Technology
Limassol, Cyprus

Martina Ziefle

Communication Science
Human–Computer Interaction
Center
RWTH Aachen University
Aachen, Germany

Introduction

A Moving Target: The Evolution of Human–Computer Interaction

Jonathan Grudin

PREAMBLE: HISTORY IN A TIME OF RAPID OBSOLESCENCE

“What is a typewriter?” my six-year-old daughter asked. I hesitated. “Well, it’s like a computer,” I began.

WHY STUDY THE HISTORY OF HUMAN–COMPUTER INTERACTION?

A paper widely read 20 years ago concluded with the advice to design a word processor by analogy to something familiar to everyone: a typewriter. Even then, one of my Danish students questioned this reading assignment noting that “the typewriter is a species on its last legs.” For most of the computing era, interaction involved 80-column punch cards, paper tape, line editors, 1920-character displays, 1-megabyte diskettes, and other extinct species. Are the interaction issues of those times relevant today? No.

Of course, aspects of the human side of human–computer interaction (HCI) change very slowly if at all. Much of what was learned about our perceptual, cognitive, social, and emotional processes when we interacted with older technologies applies to our interaction with emerging technologies as well. Aspects of how we organize and retrieve information persist, even as the specific technologies that we use change. The handbook chapters lay out relevant knowledge of human psychology; how and when that was acquired may not be critical and is not the focus here.

Nevertheless, there is a case for understanding the field’s history, and the rapid pace of change may strengthen it:

- Several disciplines are engaged in HCI research and application, but few people are exposed to more than one. By seeing how each has evolved, we can identify possible benefits of expanding our focus and obstacles to doing so.
- Celebrating the accomplishments of past visionaries and innovators is part of building a community and inspiring future contributors, even when some past achievements are difficult to appreciate today.
- Some visions and prototypes were quickly converted to widespread application, whereas others took decades and some remain unrealized to this day. By

understanding the reasons for different outcomes, we can assess today’s visions more realistically.

- Crystal balls are notoriously unreliable, but anyone planning or managing a career in a rapidly changing field must consider the future. Our best chance to anticipate change is to find trajectories that extend from the past to the present. One thing is certain: The future will not resemble the present.

This account does not emphasize engineering “firsts.” It focuses on technologies and practices as they became widely used, reflected in the spread of systems and applications. This was often paralleled by the formation of new research fields and changes in existing disciplines, which were marked by the creation and evolution of professional associations and publications. More a social history than a conceptual history, this survey points to trends and trajectories you might download into your crystal balls.

A historical account is a perspective. It emphasizes some things while de-emphasizing or omitting others. A history can be wrong in details, but is never right in any final sense. Your questions and your interests will determine how useful a perspective is to you. This introduction covers several disciplines, but the disciplines of Communication, Design, and Marketing receive less attention than another account might provide.

A blueprint for intellectual histories of HCI was established by Ron Baecker in the opening chapters of the 1987 and 1995 editions of *Readings in Human–Computer Interaction*. It was followed in Richard Pew’s chapter in the 2003 version of this handbook. Brian Shackel’s (1997) account of European contributions and specialized essays by Brad Myers (1998) on HCI engineering history and Alan Blackwell (2006) on the history of metaphor in design provide further insights and references. Perlman, Green, and Wogalter (1995) is a compendium of early HCI papers that appeared in the Human Factors literature. Research on HCI within Information Systems is covered by Banker and Kaufmann (2004) and Zhang et al. (2009). Rayward (1983, 1998) and Burke (1994, 2007) review the predigital history of information science; Burke (1998) provides a focused study of an early digital effort in this field.

In recent years many popular books covering the history of personal computing have been published (e.g., Hiltzik 1999;

Bardini 2000; Hertzfeld 2005; Markoff 2005; Moggridge 2007). This introduction extends my contribution to the previous handbook. It includes new research and draws on *Timelines* columns that have appeared in *ACM Interactions* since March 2006.

Few of the aforementioned writers are trained historians. Many lived through much of the computing era as participants and witnesses, yielding rich insights and questionable objectivity. This account draws on extensive literature and hundreds of formal interviews and discussions, but everyone has biases. Personal experiences that illustrate points can enliven an account by conveying human consequences of changes that otherwise appear abstract or distant. Some readers enjoy anecdotes, whereas others find them irritating. I try to satisfy both groups by including personal examples in a short Appendix, akin to “deleted scenes” on a DVD.

Recent years have also seen the appearance of high-quality, freely accessed digital reproductions of some early works. My references include links to several such works. The reproductions do not always preserve the original pagination, but quoted passages can be found with a search tool. Finally, all prices and costs have been converted to U.S. dollars as of 2010.

DEFINITIONS: HCI, CHI, HF&E, IT, IS, LIS

The most significant term, HCI (human–computer interaction), is defined very broadly to cover major threads of research in four disciplines: (1) Human Factors/Ergonomics (HF or HF&E), (2) Information Systems (IS), (3) Computer Science (CS), and (4) Library and Information Science (LIS). The relevant literatures are difficult to explore because they differ in the use of simple terms. This is discussed later. Here I explain how several key disciplinary labels are used. CHI (Computer-Human Interaction) has a narrower focus, associated mainly with Computer Science, the Association for Computing Machinery Special Interest Group (ACM SIGCHI), and the latter’s annual CHI conference. I use human factors and ergonomics interchangeably and refer to the discipline as HF&E—the Human Factors Society (HFS) became the Human Factors and Ergonomics Society (HFES) in 1992. (Some writers define ergonomics more narrowly around hardware.) Information Systems (IS) refers to the management discipline that has also been labeled Data Processing (DP) and Management Information Systems (MIS). I follow common parlance in referring to organizational information systems specialists as IT professionals or IT pros. With IS taken, I do not abbreviate Information Science. LIS (Library and Information Science) represents an old field with a new digital incarnation that includes important HCI research. Increasingly this discipline goes by simply “Information,” as in newly christened Schools of Information.

HUMAN–TOOL INTERACTION AND INFORMATION PROCESSING AT THE DAWN OF THE COMPUTING ERA

In the century prior to the advent of the first digital computers, advances in technology gave rise to two fields of research that later contributed to HCI: One focused on making the human

use of tools more efficient, whereas the other focused on ways to represent and distribute information more effectively.

ORIGIN OF HUMAN FACTORS

Frederick Taylor (1911) employed technologies and methods developed in the late nineteenth century—photography, moving pictures, and statistical analysis—to improve work practices by reducing performance time. Time and motion studies were applied to assembly-line manufacturing and other manual tasks. Despite the uneasiness with “Taylorism” reflected in Charlie Chaplin’s popular satire *Modern Times*, scientists and engineers strove to boost efficiency and productivity using this approach.

Lillian Gilbreth (1914) and her husband Frank were the first engineers to combine psychology and scientific management. Lillian Gilbreth focused more holistically than Taylor on efficiency and worker experience; she is regarded by some as the founder of modern Human Factors. Her PhD was the first awarded in industrial psychology. She went on to advise five U.S. presidents and became the first woman inducted into the National Academy of Engineering.

World War I and World War II accelerated efforts to match people to jobs, train them, and design equipment that could be more easily mastered. Engineering psychology was born during World War II after simple flaws in the design of aircraft controls (Roscoe 1997) and escape hatches (Dyson 1979) led to aircraft losses and thousands of casualties. Two legacies of World War II were respect for the potential of computing, based on its use in code breaking, and an enduring interest in behavioral requirements for design.

During the war, aviation engineers, psychologists, and physicians formed the Aeromedical Engineering Association. After the war, the terms “human engineering,” “human factors,” and “ergonomics” came into use, the latter primarily in Europe. For more on this history, see Roscoe (1997), Meister (1999), and HFES (2010).

Early tool use, whether by assembly-line workers or pilots, was not discretionary. If training was necessary, people were trained. One research goal was to reduce training time, but a more important goal was to increase the speed and reliability of skilled performance.

ORIGIN OF THE FOCUS ON INFORMATION

H. G. Wells, known for writing science fiction, campaigned for decades to improve society through information dissemination. In 1905, he outlined a system that might be built using another new technology of the era: index cards!

These index cards might conceivably be transparent and so contrived as to give a photographic copy promptly whenever it was needed, and they could have an attachment into which would slip a ticket bearing the name of the locality in which the individual was last reported. A little army of attendants would be at work on this index day and night. . . . An incessant stream of information would come of births, of deaths, of arrivals at inns, of applications to post offices for letters,

of tickets taken for long journeys, of criminal convictions, marriages, applications for public doles, and the like. A filter of offices would sort the stream, and all day and all night forever a swarm of clerks would go to and fro correcting this central register and photographing copies of its entries for transmission to the subordinate local stations in response to their inquiries....

Would such a human-powered “Web 2.0” be a tool for social control or public information access? The image evokes the potential, and also the challenges, of the information era that is taking shape around us now, a century later.

In the late nineteenth century, technologies and practices for compressing, distributing, and organizing information bloomed. Index cards, folders, and filing cabinets—models for icons on computer displays much later—were important inventions that influenced the management of information and organizations in the early twentieth century (Yates 1989). Typewriters and carbon paper facilitated information dissemination, as did the mimeograph machine, patented by Thomas Edison. Hollerith cards and electromechanical tabulation, celebrated steps toward computing, were heavily used to process information in industry.

Photography was used to record information as well as behavior. For almost a century, microfilm was the most efficient way to compress, duplicate, and disseminate large amounts of information. Paul Otlet, Vannevar Bush, and other microfilm advocates played a major role in shaping the future of information technology.

As the cost of paper, printing, and transportation dropped in the late nineteenth and early twentieth centuries, information dissemination and the profession of librarianship grew explosively. Library associations were formed. The Dewey Decimal and Library of Congress classification systems were developed. Thousands of relatively poorly-funded public libraries sprang up to serve local demand in the United States. In Europe, government-funded libraries were established to serve scientists and other specialists in medicine and the humanities. This difference led to different approaches to technology development on either side of the Atlantic.

In the United States, library management and the training of thousands of librarians took precedence over technology development and the needs of specialists. Public libraries adopted the simple but inflexible Dewey Decimal Classification System. The pragmatic focus of libraries and emerging library schools meant that research into technology was in the province of industry. Research into indexing, cataloging, and information retrieval was variously referred to as bibliography, documentation, and documentalism.

In contrast, the well-funded European special libraries elicited sophisticated reader demands and pressure for libraries to share resources, which promoted interest in technology and information management. The Belgian Paul Otlet obtained Melvyn Dewey’s permission to create an extended version of the Dewey Decimal System that supported what we would today call hypertext links. Otlet had to agree not to implement his “universal decimal classification” (UDC) in English for a time, an early example of a

legal constraint on technology development. UDC is still in use in some places.

In 1926, the Carnegie Foundation dropped a bombshell: It endowed the Graduate Library School (GLS) at the University of Chicago to focus solely on research. For two decades, University of Chicago was the only university granting PhDs in library studies. GLS positioned itself in the humanities and social sciences, with research into the history of publishing, typography, and other topics (Buckland 1998). *An Introduction to Library Science*, the dominant library research textbook for 40 years, was written at Chicago (Butler 1933). *It did not mention information technology at all*. Library science was shaped by the prestigious GLS program until well into the computer era, and human-tool interaction was not among its major concerns. Documentalists, researchers who focused on technology, were concentrated in industry and government agencies.

Burke (2007, p. 15) summarized the early history with its emphasis on training librarians and other specialists: “Most information professionals ... were focusing on providing information to specialists as quickly as possible. The terms used by contemporary specialists appeared to be satisfactory for many indexing tasks and there seemed no need for systems based on comprehensive and intellectually pleasing classification schemes. The goal of creating tools useful to nonspecialists was, at best, of secondary importance.”

My account emphasizes when computer technologies came into what might be called “nonspecialist use.” The early history of information management is significant, however, because the Web and declining digital storage costs have made it evident that everyone will soon become their own information managers, just as we are all now telephone operators. But I am getting ahead of our story. This section concludes with accounts of two individuals who, in different ways, shaped the history of information research and development.

Paul Otlet and the Mundaneum

Like his contemporary H.G. Wells, Otlet envisioned a vast network of information. But unlike Wells, Otlet and his collaborators built one. Otlet established a commercial research service around facts that he had been cataloging on index cards since the late nineteenth century. In 1919, the Belgian government financed the effort, which moved to a record center called the Mundaneum. By 1934, 15 million index cards and millions of images were organized using UDC, whose formula enabled the linking of items. Curtailed by the Depression and damaged during World War II, the work was largely forgotten. It was not cited by developers of the metaphorically identical Xerox NoteCards, an influential hypertext system of the 1980s.

Technological innovation continued in Europe with the development of mechanical systems of remarkable ingenuity (Buckland 2009). Features included the use of photoreceptors to detect light passing through holes in index cards positioned to represent different terms, enabling rapid retrieval of items on specific topics. These innovations inspired a well-known American scientist and research manager to go ahead with his endeavors.

Vannevar Bush and Microfilm Machines

Massachusetts Institute of Technology (MIT) professor Vannevar Bush was one of the most influential scientists in American history. He advised Presidents Franklin Roosevelt and Harry Truman, served as director of the Office of Scientific Research and Development, and was president of the Carnegie Institute.

Bush is remembered today for “As We May Think,” his 1945 *Atlantic Monthly* essay. It described the MEMEX, a hypothetical microfilm-based electromechanical information-processing machine. The MEMEX was to be a personal workstation that enabled a professional to quickly index and retrieve documents or pictures and create hypertext-like associations among them. The essay, excerpted later in this section, inspired computer engineers and computer scientists who made major contributions to HCI in the 1960s and beyond.

Not so well known is that Bush wrote the core of his essay in the early 1930s. Then, shrouded in secrecy he spent two decades and unprecedented resources on the design and construction of several machines that comprised a subset of MEMEX features. None were successful. The details are recounted in Colin Burke’s (1994) comprehensive book *Information and Secrecy: Vannevar Bush, Ultra, and the Other Memex*.

Microfilm—photographic miniaturization—had qualities that attracted Bush, as they had Otlet. Microfilm was light, could be easily transported, and was as easy to duplicate as paper records (Xerox photocopiers did not appear until 1959). The cost of handling film was brought down by technology created for the moving picture industry. Barcodelike patterns of small holes could be punched on a film and read very quickly by passing the film between light beams and photoreceptors. Microfilm was tremendously efficient as a storage medium. Memory based on relays or vacuum tubes would never be competitive, and magnetic memory, when it eventually arrived, was less versatile and far more expensive. It is easy today to overlook the compelling case that existed for basing information systems on microfilm.

Bush’s machines failed because he set overly ambitious compression and speed goals, ignored patent ownership issues, and most relevant to our account, was unaware of what librarians and documentalists had learned through decades of work on classification systems. American documentalists were active, although not well funded in their work. In 1937, the American Documentation Institute (ADI) was formed, predecessor of present-day American Society for Information Science and Technology (ASIST). Had he worked with them, Bush, an electrical engineer by training, might have avoided the fatal assumption that small sets of useful indexing terms could easily be defined and agreed upon. Metadata design is still a research challenge.

At times Bush considered libraries and the public as potential users, but his machines cost far too much for library patrons to be plausible users. He began with the Federal Bureau of Investigation (FBI) in mind and focused on military uses of cryptography and information retrieval, and a major project

was for the Central Intelligence Agency (CIA). Despite the classified nature of this work, through his academic and government positions, his writings, the vast resources he commandeered, and the scores of brilliant engineers he enlisted to work on microfilm projects, Bush promoted his vision and exerted influence for two decades, well into the computer era.

Bush’s vision emphasized both associative linking of information sources and discretionary use: Associative indexing, the basic idea of which is a provision whereby any item may be caused at will to select immediately and automatically another. This is the essential feature of the MEMEX. . . . Any item can be joined into numerous trails. . . . New forms of encyclopedias will appear, ready-made with a mesh of associative trails [which a user could extend]. . . .

The lawyer has at his touch the associated opinions and decisions of his whole experience and of the experience of friends and authorities. The patent attorney has on call the millions of issued patents, with familiar trails to every point of his client’s interest. The physician, puzzled by a patient’s reactions, strikes the trail established in studying an earlier similar case and runs rapidly through analogous case histories, with side references to the classics for the pertinent anatomy and histology. The chemist, struggling with the synthesis of an organic compound, has all the chemical literature before him in his laboratory, with trails following the analogies of compounds and side trails to their physical and chemical behavior.

The historian, with a vast chronological account of a people, parallels it with a skip trail which stops only on the salient items, and can follow at any time contemporary trails which lead him all over civilization at a particular epoch. There is a new profession of trail blazers, those who find delight in the task of establishing useful trails through the enormous mass of the common record. (Bush 1945).

Bush knew that the MEMEX was not realistic. None of his many projects included designs for the “essential” associative linking. His inspirational account nicely describes present-day hands-on discretionary use of computers by professionals. But that would arrive 50 years later, built on technologies then undreamt of. Bush did not support the early use of computers, which were slow, bulky, and expensive. Computers were clearly inferior to microfilm.

1945–1955: MANAGING VACUUM TUBES

World War II changed everything. Prior to the war, government funding of research was minimal and primarily managed by the Department of Agriculture. The unprecedented investment in science and technology during the war years revealed that huge sums could be found—for *academic or industrial research that addressed national goals*. Research expectations and strategies would never again be the same.

Sophisticated electronic computation machines built before and during World War II were designed for specific purposes, such as solving equations or breaking codes. Each of the extremely expensive cryptographic machines that helped win the war was designed to attack a specific encryption

device. A new one was needed whenever the enemy changed machines. These limitations spurred interest in general-purpose computational devices. Wartime improvements in technologies such as vacuum tubes made them more feasible, and their deployment brought HCI into the foreground.

When engineers and mathematicians emerged from military and government laboratories (and secret project rooms on university campuses), the public became aware of some of the breakthroughs. Development of ENIAC, arguably the first general-purpose computer, was begun in secret during the war but announced publicly as a “giant brain” only when it was completed in 1946. (Its first use, for calculations supporting hydrogen bomb development, was not publicized.) Accounts of the dimensions of ENIAC vary, but it stood 8–10-feet high, occupied about 1800 square feet, and consumed as much energy as a small town. It provided far less computation and memory than what can be acquired today for a few dollars, slipped into a pocket, and powered with a small battery.

Memory was inordinately expensive. Even the largest computers of the time had little memory, so they were used for computation and not for symbolic representation or information processing. Reducing operator burden was a key HCI focus, including replacing or resetting vacuum tubes more quickly, loading stored-program computers from tape rather than by manually attaching cables, and setting switches. Following “knobs and dials” human factors improvements, one computer operator could accomplish work that had previously required a team.

Libraries installed simple microfilm readers to assist the retrieval of information as publication of scholarly and popular material soared. Beyond that, library and library school involvement with technology was limited, even as the foundation for information science came into place. The war had forged alliances among the documentalists, electrical engineers, and mathematicians interested in communication and information management. Vannevar Bush’s collaborators who were involved in this effort included Claude Shannon and Warren Weaver, coauthors in 1949 of the seminal work on information theory (called communication theory at that time). Prominent American documentalist Ralph Shaw joined Bush’s efforts. Library schools continued to focus on librarianship, social science, and historical research. The GLS orientation still dominated the field. If anything the split was greater: In the 1930s, the technology-oriented ADI had included librarians and support for systems that spanned the humanities and sciences; with the coming of the war and continuing after it, ADI’s concerns became those of government and Big Science.

THREE ROLES IN EARLY COMPUTING

Early computer projects employed people in the following roles: managers, programmers, and operators. Managers oversaw the design, development, and operation of projects. They specified the programs to be written and distributed the output. Scientists and engineers wrote the programs, working with mathematically adept programmers who decomposed a

task into components that the computer could manage (for ENIAC, this was a team of six women). A small army of operators was needed. Once written, a program could take days to load by setting switches, dials, and cable connections. Despite innovations that boosted reliability, including operating vacuum tubes at lower power than normal and providing visible indicators of their failure, ENIAC was often stopped to locate and replace failed tubes. Vacuum tubes were reportedly wheeled around in shopping carts.

Eventually, each occupation—computer operation, management and systems analysis, and programming—became a major focus of HCI research, centered respectively in human factors, information systems, and computer science. Computers and our interaction with them evolved, but our research spectrum still reflects aspects of this early division of labor.

Grace Hopper: Liberating Computer Users

As computers became more reliable and capable, programming became a central activity. Computer languages, compilers, and constructs such as subroutines facilitated “programmer–computer interaction.” Grace Hopper was a pioneer in these areas. She described her goal as freeing mathematicians to do mathematics (Hopper 1952; see also Sammet 1992). This is echoed in today’s usability goal of freeing users to do their work. HCI professionals often argue that they are marginalized by software developers; in much the same way, Hopper’s accomplishments have arguably been undervalued by theoretical computer scientists.

1955–1965: TRANSISTORS, NEW VISTAS

Early forecasts that the world would need few computers reflected the limitations of vacuum tubes. Solid-state computers, which first became available commercially in 1958, changed this. Computers were still used primarily for scientific and engineering tasks, but they were reliable enough not to require a staff of computer engineers. The less computer-savvy operators who oversaw them needed better interfaces. And although computers were too expensive and limited to be widely used, the potential of transistor-based computing was evident. Some researchers envisioned possibilities that were previously unimaginable.

Another major force was reaction to the then Soviet Union’s launch of the Sputnik satellite in October 1957. This was a challenge to the West to invest in science and technology; becoming part of the response was a way to tie a research program to the national interest, which World War II had revealed to be so effective.

SUPPORTING OPERATORS: THE FIRST SYSTEMATIC HUMAN–COMPUTER INTERACTION RESEARCH

In the beginning, the computer was so costly that it had to be kept gainfully occupied for every second; people were almost slaves to feed it.

Brian Shackel (1997, p. 97)

Almost all computer use of this period involved programs and data that were read in from cards or tape. Programs then ran without interruption until they terminated, producing printed, punched, or tape output along the way. This “batch processing” restricted human interaction to basic operation, programming, and use of the output. Of these, only computer operation, the least intellectually challenging and lowest-paying job, involved hands-on computer use.

Computer operators loaded and unloaded cards and magnetic or paper tapes, set switches, pushed buttons, read lights, loaded and burst printer paper, and put printouts into distribution bins. Operators interacted directly with the system via a teletype: Typed commands interleaved with computer responses and status messages were printed on paper that scrolled up one line at a time. Eventually, they yielded to “glass tty’s” (glass teletypes), also called cathode-ray tubes (CRTs) and visual display units/terminals (VDUs/VDTs). For many years, these displays also scrolled commands and computer responses one line at a time. The price of a monochrome terminal that could display alphanumeric characters was equivalent to US\$50,000 today—expensive, but only a small fraction of the cost of the computer. A large computer might have one or more consoles. Programmers did not use the interactive consoles. Programs were typically written on paper and keypunched onto cards or tape.

Improving the design of buttons, switches, and displays was a natural extension of human factors. Experts in HF&E authored the first HCI papers. In 1959 British researcher Brian Shackel published “Ergonomics for a Computer,” followed in 1962 by “Ergonomics in the Design of a Large Digital Computer Console.” These described console redesign for analog and digital computers called the EMIac and EMIdec 2400. Shackel (1997) described the latter as the largest computer of the time.

In the United States, American aviation psychologists created the Human Engineering Society in 1956, which was focused on skilled performance including improving efficiency, reducing errors, and training. The next year it adopted the more elegant title Human Factors Society and in 1958 it initiated the journal *Human Factors*. Sid Smith’s (1963) “Man–Computer Information Transfer” marked the start of his long career with the human factors of computing.

VISIONS AND DEMONSTRATIONS

As transistors replaced vacuum tubes, a wave of imaginative writing, conceptual innovation, and prototype building swept through the research community. Some of the language is dated, notably the use of male generics, but many of the key concepts resonate even today.

J.C.R. Licklider at Bolt Beranek and Newman and Advanced Research Projects Agency

Licklider, a psychologist, played a dual role in the development of this field. He wrote influential essays and backed important research projects as a manager at Bolt Beranek and Newman (BBN) from 1957 to 1962 and as director of

the Information-Processing Techniques Office (IPTO) of the Department of Defense Advanced Research Projects Agency (called ARPA and DARPA at different times) from 1962 to 1964.

BBN employed dozens of influential researchers on computer-related projects funded by the government, including John Seely Brown, Richard Pew, and many MIT faculty members such as John McCarthy, Marvin Minsky, and Licklider himself. Funding by IPTO was crucial in creating computer science departments and establishing artificial intelligence (AI) as a discipline in the 1960s. It is best known for a Licklider project that created the forerunner of the Internet called the ARPANET.

In 1960, Licklider outlined a vision he called *man–machine symbiosis*: “There are many man–machine systems. At present, however, there are no man–computer symbioses—answers are needed.” The computer was “a fast information-retrieval and data-processing machine” destined to play a larger role: “One of the main aims of man–computer symbiosis is to bring the computing machine effectively into the formulative parts of technical problems” (pp. 4–5).

This required rapid, real-time interaction, which batch systems did not support. In 1962, Licklider and Wes Clark outlined the requirements of a system for “online man–computer communication.” They identified capabilities that they felt were ripe for development: time-sharing of a computer among many users; electronic input–output surfaces to display and communicate symbolic and pictorial information; interactive, real-time support for programming and information processing; large-scale information storage and retrieval systems; and facilitation of human cooperation. They foresaw that other desirable technologies, such as speech recognition and natural language understanding, would be very difficult to achieve.

In a 1963 memorandum that cleverly tied computing to the emerging post-Sputnik space program, Licklider addressed his colleagues as “the members and affiliates of the Intergalactic Computer Network” and identified many features of a future Internet (Licklider 1963). His 1965 book *Libraries of the Future* expanded this vision. Licklider’s role in advancing computer science and HCI is detailed by Waldrop (2001).

John McCarthy, Christopher Strachey, and Wesley Clark

McCarthy and Strachey worked out details of time-sharing, which made interactive computing possible (Fano and Corbato 1966). Apart from a few researchers who had access to computers built with no-expenses-spared military funding, computer use was too expensive to support exclusive individual access. Time-sharing allowed several (and later dozens) simultaneous users to work at terminals. Languages were developed to facilitate the control and programming of time-sharing systems (e.g., JOSS in 1964).

Clark was instrumental in building the TX-0 and TX-2 at MIT’s Lincoln Laboratory to demonstrate time-sharing and other innovative concepts. These machines, which cost on the order of US\$10 million, helped establish the Boston area as a center for computer research. The TX-2 was the

most powerful and capable computer in the world at the time. It was much less powerful and capable than a present-day smartphone. Clark and Ivan Sutherland discussed this era in a CHI'05 panel, which is accessible online (Buxton 2006).

Ivan Sutherland and Computer Graphics

Sutherland's 1963 PhD thesis may be the most influential document in the history of HCI. His Sketchpad system, built on TX-2 to make computers "more approachable," launched computer graphics, which would have a decisive impact on HCI 20 years later. A nice version restored by Alan Blackwell and Kerry Rodden is available (<http://www.cl.cam.ac.uk/TechReports/UCAM-CL-TR-574.pdf>).

Sutherland demonstrated iconic representations of software constraints, object-oriented programming concepts, and the copying, moving, and deleting of hierarchically organized objects. He explored novel interaction techniques, such as picture construction using a light pen. He facilitated visualization by separating the coordinate system used to define a picture from the one used to display it, and demonstrated animated graphics, noting the potential for digitally rendered cartoons 20 years before *Toy Story*. His frank descriptions enabled others to make rapid progress in the field—when engineers found Sketchpad too limited for computer-assisted design (CAD), he called the trial a "big flop" and indicated why.

In 1964, with his PhD behind him, Sutherland succeeded Licklider as the director of IPTO. Among those he funded was Douglas Engelbart at the Stanford Research Institute (SRI).

Douglas Engelbart: Augmenting Human Intellect

In 1962, Engelbart published "Augmenting Human Intellect: A Conceptual Framework." Over the next several years he built systems that made astonishing strides toward realizing this vision. He also supported and inspired engineers and programmers who went on to make major independent contributions.

Echoing Bush and Licklider, Engelbart saw the potential for computers to become congenial tools that people would choose to use interactively:

By 'augmenting human intellect' we mean increasing the capability of a man to approach a complex problem situation, to gain comprehension to suit his particular needs, and to derive solutions to problems. ... By 'complex situations' we include the professional problems of diplomats, executives, social scientists, life scientists, physical scientists, attorneys, designers. ... We refer to a way of life in an integrated domain where hunches, cut-and-try, intangibles, and the human 'feel for a situation' usefully coexist with powerful concepts, streamlined terminology and notation, sophisticated methods, and high-powered electronic aids.

(Engelbart 1962, p. 1)

Engelbart used ARPA funding to rapidly develop and integrate an extraordinary set of prototype applications into his NLS system. In doing so, he conceptualized and implemented the foundations of word processing, invented or refined input devices including the mouse and the multikey control box, and made use of multidisplay environments that

integrated text, graphics, and video in windows. These unparalleled advances were demonstrated in a sensational 90-minute live event at the 1968 Fall Joint Computer Conference in San Francisco, California (<http://sloan.stanford.edu/MouseSite/1968Demo.html>). The focal point for interactive systems research in the United States was moving from the East Coast to the West Coast.

Engelbart, an engineer, supported human factors testing to improve efficiency and reduce errors in skilled use, focusing on effects of fatigue and stress. Engelbart's systems required training. He felt that people should be willing to tackle a difficult interface if it delivered great power once mastered. Unfortunately, the lack of concern for initial usability was a factor in Engelbart's loss of funding. His demonstration became something of a success disaster: DARPA was impressed and installed NLS, but found it too difficult to use (Bardini 2000). Years later, the question "Is it more important to optimize for skilled use or initial use?" was widely debated, and still occasionally surfaces in HCI discussions.

Ted Nelson's Vision of Interconnectedness

In 1960, Ted Nelson, a graduate student in sociology who coined the term hypertext, founded Project Xanadu. The goal was an easily used computer network. In 1965, he published a paper titled "A File Structure for the Complex, the Changing and the Indeterminate." Nelson continued to write stirring calls for systems to democratize computing through a highly interconnected, extensible network of digital objects (e.g., Nelson 1973). Xanadu was never fully realized. Nelson did not consider the early World Wide Web to be an adequate realization of his vision, but lightweight technologies such as weblogs, wikis, collaborative tagging, and search enable many of the activities he envisioned.

Later, Nelson (1996) foresaw intellectual property issues arising in digital domains and coined the term "micropayment." Although his solutions were again not fully implemented, they drew attention to important issues.

FROM DOCUMENTATION TO INFORMATION SCIENCE

The late 1950s saw the last major investments in microfilm and other predigital systems. The most ambitious were military and intelligence systems, including Vannevar Bush's final efforts (Burke 1994). Documentalists began to see that declining memory costs would enable computation engines to become information-processing machines. The conceptual evolution was relatively continuous, but at the institutional level change could come swiftly. New professions—mathematicians and engineers—were engaged in technology development, new initiatives were launched that still bore few ties to contemporary librarianship or the humanities orientation of library schools. A new banner was needed.

Merriam Webster dates the term information science to 1960. Conferences held at Georgia Institute of Technology in 1961 are credited with shifting the focus from information as a technology to information as an incipient science. In 1963, chemist-turned-documentalist Jason Farradane taught the

first information science courses at City University, London, United Kingdom. The profession of chemistry had long invested in organizing its literature systematically, and another chemist-turned-documentalist Allen Kent was at the center of a major information science initiative at the University of Pittsburgh (Aspray 1999). In the early 1960s, Anthony Debons, a psychologist and friend of Licklider, organized a series of NATO-sponsored congresses at Pittsburgh. Guided by Douglas Engelbart, these meetings centered on people and on how technology could augment their activities. In 1964 the Graduate Library School at the University of Pittsburgh became the Graduate School of Library and Information Sciences, and Georgia Tech formed a School of Information Science initially with one full-time faculty member.

CONCLUSION: VISIONS, DEMOS, AND WIDESPREAD USE

Progress in HCI can be understood in terms of inspiring visions, conceptual advances that enable aspects of the visions to be demonstrated in working prototypes, and the evolution of design and application. The engine, enabling visions to be realized and soon thereafter to be widely deployed, was the relentless hardware advance that produced devices that were millions of times more powerful than the much more expensive systems designed and used by the pioneers.

At the conceptual level, much of the basic foundation for today's graphical user interfaces (GUIs) was in place by 1965. However, at that time it required individual use of a US\$10-million custom-built machine. Pew (2003, p. 3) describes the 1960 Digital Equipment Corporation (DEC) PDP-1 as a breakthrough, "truly a computer with which an individual could interact." The PDP-1 came with a CRT display, keyboard, light pen, and paper tape reader. It cost about US\$1 million and had the capacity that a Radio Shack TRS 80 had 20 years later. It required considerable technical and programming support. Even the PDP-1 could only be used by a few fortunate researchers.

Licklider's man-computer symbiosis, Engelbart's augmenting human intellect, and Nelson's "conceptual framework for man-machine everything" described a world that did not exist. It was a world in which attorneys, doctors, chemists, and designers chose to become hands-on users of computers. For some time to come, the reality would be that most hands-on users were computer operators engaged in routine, nondiscretionary tasks. As for the visions, 40 years later some of the capabilities are taken for granted, some are just being realized, and others remain elusive.

1965–1980: HUMAN-COMPUTER INTERACTION PRIOR TO PERSONAL COMPUTING

Control Data Corporation launched the transistor-based 6000 series computer in 1964. In 1965, commercial computers based on integrated circuits arrived with the IBM System/360. These powerful systems, later called mainframes to distinguish them from minicomputers, firmly established computing in the business realm. Each of the three computing

roles—operation, management, and programming—became a significant profession.

Operators still interacted directly with computers for routine maintenance and operation, and as time-sharing developed, hands-on use expanded to include data entry and other repetitive tasks. Managers and systems analysts oversaw hardware acquisition, software development, operation, and the use of output. They were usually not hands-on users, although people who relied on printed output and reports did call themselves "computer users."

Apart from those working in research settings, few programmers were direct users until late in this period. Many prepared flowcharts and wrote programs on paper forms. Key punch operators then punched the program instructions onto cards, which were sent to computer centers for computer operators to load into the computer and run. Printouts and other output were picked up later. Many programmers used computers directly when they could, but the cost generally dictated more efficient division of labor.

We are focusing on broad trends. Business computing took off in the mid-1960s, although the 1951 LEO I was probably the first commercial business computer. This interesting venture, which ended with the arrival of the mainframe era, is detailed in Wikipedia (under 'LEO computer') and the books and articles referenced there.

HUMAN FACTORS AND ERGONOMICS EMBRACE COMPUTER OPERATION

In 1970, Brian Shackel founded the Human Sciences and Advanced Technology (HUSAT) center at Loughborough University in Leicestershire, the United Kingdom, which is devoted to ergonomics research that emphasizes HCI. Sid Smith and other human factors engineers worked on input and output issues, such as the representation of information on displays (e.g., Smith, Farquhar, and Thomas 1965) and computer-generated speech (Smith and Goodwin 1970). The Computer Systems Technical Group (CSTG) of the HFS was formed in 1972, and soon it was the largest technical group in the society.

The general *Human Factors* journal was joined in 1969 by the computer-focused *International Journal of Man-Machine Studies (IJMMS)*. The first widely read HCI book was James Martin's (1973) *Design of Man-Computer Dialogues*. Martin's comprehensive survey of interfaces for operation and data entry began with an arresting opening chapter that described a world in transition. Extrapolating from declining hardware prices, he wrote, "The terminal or console operator, instead of being a peripheral consideration, will become the tail that wags the whole dog. ... The computer industry will be forced to become increasingly concerned with the usage of people, rather than with the computer's intestines" (pp. 3–4).

In the mid-1970s, U.S. government agencies responsible for agriculture and social security initiated large-scale data-processing system projects, described by Pew (2003). Although not successful, these efforts led to methodological innovations in the use of style guides, usability laboratories, prototyping, and task analysis.

In 1980, three significant HF&E books were published: two on VDT design (Cakir, Hart, and Stewart 1980; Grandjean and Vigliani 1980) and one general guideline (Damodaran, Simpson, and Wilson 1980). Drafts of a German work on VDT standards, made public in 1981, provided an economic incentive to design for human capabilities by threatening to ban noncompliant products. Later in the same year, a corresponding American National Standards Institute standards group for “office and text systems” was formed.

INFORMATION SYSTEMS (IS) ADDRESSES THE MANAGEMENT OF COMPUTING

Companies acquired expensive business computers to address major organizational concerns. Even when the principal concern was simply to appear modern (Greenbaum 1979), the desire to show benefits from a multimillion dollar investment could chain managers to a computer almost as tightly as were the operator and data entry “slaves.” In addition to being expected to make use of output, they might encounter resistance to system acceptance.

Beginning in 1967, the journal *Management Science* published a column titled “Information Systems in Management Science.” Early definitions of IS included “an integrated man-machine system for providing information to support the operation, management, and decision-making functions in an organization” (Davis 1974) and “the effective design, delivery, and use of information systems in organizations” (Keen 1980 quoted in Zhang, Nah, and Preece 2004). In 1968, an MIS center and degree program was established at Minnesota. It initiated several influential research streams and in 1977 launched *MIS Quarterly*, the leading journal in the field. The MIS field juxtaposed a focus on specific tasks in organizational settings with demands for general theory and precise measurement, a challenging combination.

A historical survey (Banker and Kaufmann 2004) identifies HCI as one of five major IS research streams and dates it back to Ackoff’s (1967) paper describing challenges in handling computer-generated information. There was some research into hands-on operator issues such as data entry and error messages, but for a decade most HCI work in IS dealt with the users of information, typically managers. Research included the design of printed reports, but the drive for theory led to a strong focus on cognitive styles: individual differences in how people (notably managers) perceive and process information. Articles on HCI were published in the human factors-oriented *IJMM*s as well as management journals.

Sociotechnical approaches to system design (Mumford 1971, 1976; Bjørn-Andersen and Hedberg 1977) were developed in response to user difficulties and resistance. These involved educating representative workers about technological possibilities and involving them in design, in part to increase their acceptance of the resulting system. Late in this period, sophisticated views of the complex social and organizational dynamics around system adoption and use emerged (e.g., Kling 1980; Markus 1983).

PROGRAMMING: SUBJECT OF STUDY, SOURCE OF CHANGE

Even programmers who were not hands-on users were interacting with computers, and more than 1000 research papers on variables affecting programming performance were published in the 1960s and 1970s (Baecker and Buxton 1987). Most were studies of the behavior of programmers in isolation, independent of organizational context. Influential reviews of this work included Gerald Weinberg’s landmark *The Psychology of Computer Programming* in 1971; Ben Shneiderman’s *Software Psychology: Human Factors in Computer and Information Systems* in 1980; and Beau Sheil’s 1981 review of studies of programming notation (conditionals, control flow, data types), practices (flowcharting, indenting, variable naming, commenting), and tasks (learning, coding, debugging).

Software developers changed the field through invention. In 1970, Xerox Palo Alto Research Center (PARC) was founded to advance computer technology by developing new hardware, programming languages, and programming environments. It attracted researchers and system builders from the laboratories of Engelbart and Sutherland. In 1971, Allen Newell of Carnegie Mellon University (CMU), Pennsylvania, proposed a project to PARC, which was launched 3 years later: “Central to the activities of computing—programming, debugging, etc.—are tasks that appear to be within the scope of this emerging theory [a psychology of cognitive behavior]” (Card and Moran 1986, p. 183).

Like HUSAT, which was also launched in 1970, PARC had a broad charter. HUSAT focused on ergonomics, anchored in the tradition of nondiscretionary use, one component of which was the human factors of computing. PARC focused on computing, anchored in visions of discretionary use, one component of which was also the human factors of computing. Researchers at PARC, influenced by cognitive psychology, extended the primarily perceptual motor focus of human factors to higher-level cognition, whereas HUSAT, influenced by sociotechnical design, extended human factors by considering organizational factors.

COMPUTER SCIENCE: A NEW DISCIPLINE

Computer science departments in educational institutions emerged in the mid-1960s. Some originated in engineering, others in applied mathematics. From engineering, computer graphics was a specialization of particular relevance to HCI. Applied mathematics was the background of many early AI researchers, which has interacted with HCI in complex ways in subsequent years.

The expensive early machines capable of interesting work were funded without consideration to cost by branches of the military. Technical success was the sole evaluation criterion (Norberg and O’Neill 1996). Directed by Licklider, Sutherland, and their successors, ARPA played a major role. The need for heavy funding concentrated researchers in a few centers, which bore little resemblance to the batch and time-shared business computing environments of that era.

User needs differed: The technically savvy hands-on users in research settings did not press for low-level interface enhancements.

The computer graphics and AI perspectives that arose in these centers differed from the perspectives of HCI researchers who focused on less expensive, more widely deployed systems. Computer graphics and AI required processing power; hardware advances meant declining cost for the same high level of computation. For HCI researchers, hardware advances meant greater computing capability at the same low price. Only later would this difference diminish, when widely available machines could support graphical interfaces and some AI programs. Despite this gap, between 1965 and 1980 some computer science researchers focused on interaction, which is not surprising given that interaction was an element of the visions formulated in the previous decade.

Computer Graphics: Realism and Interaction

In 1968, Sutherland joined David Evans to establish an influential computer graphics laboratory at the University of Utah. The Utah Computer Science Department was founded in 1965, as part of computer science's first move into academic prominence. Utah contributed to the western migration as graduates of the laboratory, including Alan Kay and William Newman (and later Jim Blinn and Jim Clark), went to California. Most graphics systems at the time were built on the DEC PDP-1 and PDP-7. These expensive machines—the list price of a high-resolution display alone was equivalent to more than US\$100,000 in today's dollars—were in principle capable of multitasking, but in practice most graphics programs required all of a processor's cycles.

In 1973 the Xerox Alto arrived, a powerful step toward realizing Alan Kay's vision of computation as a medium for personal computing (Kay and Goldberg 1977). The Alto was too expensive to be widely used—it was never widely marketed—and not powerful enough to support high-end graphics research, but it did support graphical interfaces of the kind Engelbart had prototyped. In doing so, the Alto signaled the approach of inexpensive, interactive, personal machines capable of supporting graphics. Computer graphics researchers had to decide whether to focus on high-end graphics or on more primitive features that would soon run on widely affordable machines.

William Newman, coauthor in 1973 of the influential *Principles of Interactive Computer Graphics*, described the shift in a personal communication: "Everything changed—the computer graphics community got interested in realism; I remained interested in interaction, and I eventually found myself doing HCI." He was not alone. Other graphics researchers whose focus shifted to broader interaction issues included Ron Baecker and Jim Foley. Foley and Wallace (1974, p. 462) identified requirements for designing "interactive graphics systems whose aim is good symbiosis between man and machine." The shift was gradual: A total of 18 papers in the first SIGGRAPH conference, in 1974, had the words "interactive" or "interaction" in their titles. A decade later, there would be none.

At Xerox, Larry Tesler and Tim Mott recognized that Alto could support a graphical interface accessible to untrained people. The latter point had not been important given the prior focus on trained, expert performance. By early 1974, Tesler and Mott had developed the Gypsy text editor. Gypsy and Xerox's Bravo editor developed by Charles Simonyi preceded and influenced Microsoft Word (Hiltzik 1999).

The focus on interaction was highlighted in 1976 when SIGGRAPH sponsored a 2-day workshop in Pittsburgh, User-Oriented Design of Interactive Graphics Systems (UODIGS). Participants who were later active in CHI included Jim Foley, William Newman, Ron Baecker, John Bennett, Phyllis Reisner, and Tom Moran. Licklider and Nicholas Negroponte presented vision papers. The conference was managed by the chair of Pittsburgh's computer science department. One participant was Anthony Debons, Licklider's friend who had helped build Pittsburgh's world-renowned information science program. The UODIGS'76 workshop arguably marked the end of a visionary period, embodying an idea whose time had not quite yet come. Licklider saw it clearly:

Interactive computer graphics appears likely to be one of the main forces that will bring computers directly into the lives of very large numbers of people during the next two or three decades. Truly user-oriented graphics of sufficient power to be useful to large numbers of people has not been widely affordable, but it will soon become so and, when it does, the appropriateness and quality of the products offered will to a large extent determine the future of computers as intellectual aids and partners of people.

(Licklider 1976, p. 89)

UODIGS was not repeated. Despite the stature of its participants, the 150-page proceedings were not cited. Not until 1981 was another user-oriented design conference held, after which such conferences were held every year. Application of graphics was not quite at hand; most HCI research remained focused on interaction driven by commands, forms, and full-page menus.

Artificial Intelligence: Winter Follows Summer

In the late 1960s and early 1970s AI burst onto the scene, promising to transform HCI. It did not go as planned. Logically, AI and HCI are closely related. What are intelligent machines for if not to interact with people? Research on AI has influenced HCI: Speech recognition and natural language are perennial HCI topics; expert, knowledge-based, adaptive, and mixed-initiative systems have been tried, as have applications of production systems, neural networks, and fuzzy logic. Today, human–robot interaction and machine learning are attracting much attention.

Although some AI features make it into systems and applications, frequent predictions that powerful machines would soon bring major AI technologies into wide use and thus become a focus of HCI research were not borne out. AI did not come into focus in HCI, and AI researchers showed limited interest in HCI.

To piece this together one requires a brief review of early AI history. The term “artificial intelligence” first appeared in a 1955 call by John McCarthy for a meeting on machine intelligence that was held in Dartmouth. In 1956, Alan Turing’s prescient essay “Computing Machinery and Intelligence” attracted attention when it was reprinted in *The World of Mathematics*. (It was first published in 1950, as were Claude Shannon’s “Programming a Computer for Playing Chess” and Isaac Asimov’s *I, Robot*, which explored his three laws of robotics.) Newell and Simon presented their logic theory machine in 1956 and then focused on developing a general problem solver. McCarthy invented the LISP programming language in 1958 (McCarthy 1960).

Many AI pioneers were trained in mathematics and logic, where almost everything can be derived from a few axioms and a small set of rules. Mathematical ability is considered a high form of intelligence, even by non-mathematicians. AI researchers anticipated that machines that operate logically and tirelessly would achieve high levels of intelligence—applying a small set of rules to a limited number of objects. Early AI focused on theorem-proving and games and problems that had a strong logical focus, such as chess and go. McCarthy (1988), who espoused predicate calculus as a foundation for AI, summed it up as follows:

As suggested by the term ‘artificial intelligence’, we were not considering human behavior except as a clue to possible effective ways of doing tasks. The only participants who studied human behavior were Newell and Simon. (The goal) was to get away from studying human behavior and consider the computer as a tool for solving certain classes of problems. Thus, AI was created as a branch of computer science and not as a branch of psychology.

Unfortunately, by ignoring psychology, mathematicians overlooked the complexity and inconsistency that mark human beings and our social constructs. Underestimating the complexity of intelligence, they overestimated the prospects for creating it artificially. Hyperbolic predictions and AI have been close companions. In the summer of 1949 the British logician and code breaker Alan Turing wrote in the *London Times*:

I do not see why [the computer] should not enter any one of the fields normally covered by the human intellect, and eventually compete on equal terms. I do not think you can even draw the line about sonnets, though the comparison is perhaps a little bit unfair because a sonnet written by a machine will be better appreciated by another machine.

Optimistic forecasts by the 1956 Dartmouth workshop participants attracted considerable attention. When they collided with reality, a pattern was established that was to play out repeatedly. Hans Moravec (1998) wrote:

In the 1950s, the pioneers of AI viewed computers as locomotives of thought, which might outperform humans in higher mental work as prodigiously as they outperformed them in

arithmetic, if they were harnessed to the right programs. . . .

By 1960 the unspectacular performance of the first reasoning and translation programs had taken the bloom off the rose.

A significant part of the pattern is that HCI thrives on resources that are freed when interest in AI declines. In 1960, with the bloom wearing off the AI rose, the managers of MIT’s Lincoln Laboratory looked for new uses for the massive government-funded TX-0 and TX-2 computers. Ivan Sutherland’s Sketchpad and early computer graphics were a result.

The response to Sputnik reversed the downturn in AI prospects. Licklider, as director of ARPA’s IPTO (1962–1964), provided extensive support for computer science in general and AI in particular. MIT’s Project Mac, founded in 1963 by Marvin Minsky and others, initially received US\$13 million per year, rising to US\$24 million in 1969. ARPA sponsored the AI Laboratory at SRI, AI research at CMU, and Nicholas Negroponte’s Machine Architecture Group at MIT. A dramatic early achievement, SRI’s Shakey the Robot, was featured in articles in *Life* (Darrach 1970) and *National Geographic* (White 1970). Given a simple but nontrivial task, Shakey could apparently go to the desired location, scan and reason about the surroundings, and move objects as needed to accomplish the goal (for Shakey at work, see <http://www.ai.sri.com/shakey/>).

In 1970, Negroponte outlined a case for machine intelligence: “Why ask a machine to learn, to understand, to associate courses with goals, to be self-improving, to be ethical—in short, to be intelligent?” He noted common reservations, “People generally distrust the concept of machines that approach (and thus why not pass?) our own human intelligence,” and identified a key problem: “Any design procedure, set of rules, or truism is tenuous, if not subversive, when used out of context or regardless of context.” This insight, that it is risky to apply algorithms without understanding the situation at hand, led Negroponte to a false inference: “*It follows that a mechanism must recognize and understand the context before carrying out an operation.*” (Negroponte 1970, p. 1; my italics).

A perfectly reasonable alternative is that the mechanism is guided by humans who understand the context: Licklider’s human–machine symbiosis. Overlooking this, Negroponte built a case for an ambitious research program:

Therefore, a machine must be able to discern changes in meaning brought about by changes in context, hence, be intelligent. And to do this, it must have a sophisticated set of sensors, effectors, and processors to view the real world directly and indirectly. . . . A paradigm for fruitful conversations must be machines that can speak and respond to a natural language. . . . But, the *tête-à-tête* [*sic*] must be even more direct and fluid; it is gestures, smiles, and frowns that turn a conversation into a dialogue. . . . Hand waving often carries as much meaning as text. Manner carries cultural information: The Arabs use their noses, the Japanese nod their heads. . . . Imagine a machine that can follow your design methodology and at the same time discern and assimilate your conversational idiosyncrasies. This same machine after observing

your behavior could build a predictive model of your conversational performance. Such a machine could then reinforce the dialogue by using the predictive model to respond to you in a manner that is in rhythm with your personal behavior and conversational idiosyncrasies. . . . The dialogue would be so intimate—even exclusive—that only mutual persuasion and compromise would bring about ideas, ideas unrealizable by either conversant alone. No doubt in such a symbiosis it would not be solely the human designer who would decide when the machine is relevant (pp. 1–13).

The same year, Negroponte's MIT colleague Minsky went further, as reported in *Life*:

In from three to eight years we will have a machine with the general intelligence of an average human being. I mean a machine that will be able to read Shakespeare, grease a car, play office politics, tell a joke, and have a fight. At that point, the machine will begin to educate itself with fantastic speed. In a few months, it will be at genius level and a few months after that its powers will be incalculable.

(Darrach 1970, p. 60)

Other AI researchers told Darrach that Minsky's timetable was ambitious: "Give us 15 years was a common remark—but all agreed that there would be such a machine and that it would precipitate the third Industrial Revolution; wipe out war and poverty; and roll up centuries of growth in science, education, and the arts" (Darrach 1970, p. 60).

Such predictions were common. In 1960, Nobel laureate and AI pioneer Herb Simon wrote: "Machines will be capable, within 20 years, of doing any work that a man can do." (Simon 1960, p. 38). Five years later, I. J. Good, an Oxford mathematician, wrote, "The survival of man depends on the early construction of an ultraintelligent machine" that "could design even better machines; there would then unquestionably be an 'intelligence explosion', and the intelligence of man would be left far behind" (Good 1965, pp. 31–33).

The Darrach article ended by quoting Ross Quillian:

I hope that man and these ultimate machines will be able to collaborate without conflict. But if they can't, we may be forced to choose sides. And if it comes to choice, I know what mine will be. My loyalties go to intelligent life, no matter in what medium it may arise".

(Darrach 1970, p. 68)

It is important to understand the anxieties of the time and the consequences of such claims. The world had barely avoided a devastating thermonuclear war during the Cuban missile crisis of 1962. Leaders seemed powerless to defuse the Cold War. Responding to a sense of urgency, ARPA initiated major programs in speech recognition and natural language understanding in 1971.

Ironically, central to funding this research was a psychologist not wholly convinced by the vision. Citing an Air Force study that predicted that intelligent machines might take 20 years to arrive, Licklider (1960) noted that in this interval HCI would be useful: "That would leave, say, 5 years to

develop man–computer symbiosis and 15 years to use it. The 15 may be 10 or 500, but those years should be intellectually the most creative and exciting in the history of mankind." Ten to five hundred years represent breathtaking uncertainty. Recipients of Licklider's funding were on the optimistic end of this spectrum.

Five years later, disappointed with the progress, ARPA discontinued speech and language support—for a while. In Europe, a similar story unfolded. Through the 1960s, AI research expanded in Great Britain. A principal proponent was Turing's former colleague Donald Michie. Then in 1973 the Lighthill report, commissioned by the Science and Engineering Research Council, reached generally negative conclusions about AI's prospects for scaling up to address real-world problems. Almost all government funding was cut off.

The next decade was an AI winter, a recurring season in which research funding is withheld due to disillusionment over unfulfilled promises. The bloom was again off the rose, but it would prove to be a hardy perennial (Grudin 2009).

LIBRARY SCHOOLS EMBRACE INFORMATION SCIENCE

Early information science research and studies of "human information behavior" were initiated in the 1960s and 1970s, which focused on scholarship and application in science and engineering (Fidel 2011). The response to Sputnik proved that Big Science research did not end when the war ended. Aligning their work with national priorities became a priority for many researchers.

The terms "information science," "information technology," and "information explosion" swept into use. The Pittsburgh and Georgia Tech programs flourished. Pittsburgh created the first information science PhD program in the United States in 1970, identifying humans "as the central factor in the development of an understanding of information phenomena" (Aspray 1999, p. 12). The program balanced behavioral sciences (psychology, linguistics, communication) and technical grounding (automata theory, computer science). In 1973, Pittsburgh established the first information science department. Its program developed a strong international reputation. Slowly, the emphasis shifted from behavior to technology. On being awarded a major National Science Foundation (NSF) center grant in 1966, the Georgia Tech school expanded. In 1970 it became a PhD-granting school, rechristened as Information and Computer Science.

In 1968, the American Documentation Institute became the American Society for Information Science, and 2 years later the journal *American Documentation* became *Journal of the American Society for Information Science*. In 1978, the ACM Special Interest Group on Information Retrieval (SIGIR) was formed. It launched an annual conference for "Information Storage and Retrieval" (since 1982, "Information Retrieval"), modeled on a 1971 conference. In 1984, the American Library Association belatedly embraced the i-word by creating the Association for Library and Information Science Education (ALISE), which convened an annual research conference.

By 1980, schools at over a dozen universities had added the word information to their titles. Many were library school transitions. Delivery on the promise of transformative technology lagged, however. For example, from 1965 to 1972 the Ford and Carnegie Foundations, NSF, DARPA, and the American Newspaper Publishers Association invested over US\$30 million in MIT's Project Intrex (Burke 1998). The largest nonmilitary information research project of its time, Intrex was to be the library of the future. Online catalogs were to include up to 50 index fields per item, accessible on CRT displays, with full text of books and articles converted to microfilm and read via television displays. None of this proved feasible.

Terminal-based computing costs declined. The ARPANET debuted in 1969, and supported e-mail in 1971 and file sharing in 1973. This spurred visions of a "network society" of the future (Hiltz and Turoff 1978).

As an aside, the technological optimism that marked this era lacked the nuanced psychological insight of E. M. Forster who in 1909 anticipated AI and networking developments in his remarkable story *The Machine Stops*.

1980–1985: DISCRETIONARY USE COMES INTO FOCUS

In 1980, most HF&E and IS research focused on the down-to-earth business of making efficient use of expensive mainframes. The beginning of a major shift went almost unnoticed. Less expensive but highly capable minicomputers based on LSI technology enabled DEC, Wang Laboratories, and Data General to make inroads into the mainframe market. At the low end, home computers gained traction. Students and hobbyists were drawn to these minis and micros, creating a population of hands-on discretionary users. There were experimental trials of online library catalogs and electronic journals.

Then, between 1981 and 1984 a flood of innovative and powerful computers were released: Xerox Star; IBM PC; Apple Lisa; LISP machines from Symbolics and Lisp Machines, Inc. (LMI); workstations from Sun Microsystems and Silicon Graphics; and the Apple Macintosh. On January 1, 1984, AT&T's breakup into competing companies took effect. AT&T had more employees and more customers than any other U.S. company. It was a monopoly: Neither its customers nor its employees had discretion in technology use. Both AT&T and its Bell Laboratories research division had employed human factors research to improve training and increase efficiency. Suddenly freed from a ban on entering the computer business, AT&T launched the ill-fated Unix PC in 1985. AT&T and the new regional operating companies now faced customers who had choices, and their HCI focus broadened accordingly (Israelski and Lund 2003).

In general, lower-priced computers created markets for shrink-wrap software. For the first time, computer and software companies targeted significant numbers of nontechnical hands-on users who received little or no formal training.

It had taken 20 years, but early visions were being realized. Nonprogrammers were choosing to use computers to do their work. The psychology of discretionary users intrigued two groups: (1) psychologists who liked to use computers and (2) technology companies who wanted to sell to discretionary users. Not surprisingly, computer and telecommunication companies started hiring a lot of experimental psychologists.

DISCRETION IN COMPUTER USE

Technology use lies on a continuum bracketed by the assembly-line nightmare of *Modern Times* and the utopian vision of completely empowered individuals. To use a technology or not to use it—sometimes we have a choice, other times we do not. On the phone, we may have to wrestle with speech recognition and routing systems. At home, computer use may be largely discretionary. The workplace often lies in between: Technologies are prescribed or proscribed, but we ignore some injunctions or obtain exceptions, we use some features but not others, and we join with colleagues to press for changes.

For early computer builders, work was more a calling than a job, but operation required a staff to carry out essential if less interesting tasks. For the first half of the computing era, most hands-on use was by people with a mandate. Hardware innovation, more versatile software, and steady progress in understanding the psychology of users and tasks—and transferring that understanding to software developers—led to hands-on users who had more choice regarding how they worked. Rising expectations played a role; people learned that software is flexible and expected it to be more congenial. Competition among vendors produced alternatives. With more emphasis on marketing to consumers came more emphasis on user-friendliness.

Discretion is not all-or-none. No one must use a computer, but many jobs and pastimes require it. People can resist, sabotage, or quit their jobs. However, a clerk or a systems administrator has less discretion than someone using technology for a leisure activity. For an airline reservation clerk, computer use is mandatory. For a traveler booking a flight, computer use is discretionary. This distinction, and the shift toward greater discretion, is at the heart of the history of HCI.

The shift was gradual. About 30 years ago, John Bennett (1979) predicted that discretionary use would lead to more emphasis on usability. The 1980 book *Human Interaction with Computers*, edited by Harold Smith and Thomas Green, perched on the cusp. It included an article by Jens Rasmussen, "The Human As a Systems Component," that covered the nondiscretionary perspective. One-third of the book covered research on programming. The remainder addressed "non-specialist people," discretionary users who are not computer savvy. Smith and Green wrote, "It is not enough just to establish what computer systems can and cannot do; we need to spend just as much effort establishing what people can *and want to do*" (p. viii, italics in original).

A decade later, Liam Bannon (1991) noted broader implications of a shift "from human factors to human actors." The

trajectory is not always toward choice. Discretion can be curtailed—for example, word processor use is now often a job requirement and not an alternative to using a typewriter. Even in an era of specialization, customization, and competition, the exercise of choice varies over time and across contexts. Discretion is only one factor, but an analysis of its role casts light on how HCI efforts differ and why they have remained distinct through the years.

MINICOMPUTERS AND OFFICE AUTOMATION

Cabinet-sized minicomputers that could support several people were available from the mid-1960s. By late 1970s, superminis such as the VAX 11/780 supported integrated suites of productivity tools. In 1980, DEC, Data General, and Wang Laboratories were growth companies near Boston.

A minicomputer could handle personal productivity tools or a database of moderate size. Users sat at terminals. With “dumb terminals,” the central processor handled each keystroke. Other terminals had a processor that supported a user who entered a screenful of data, which was then on command sent as a batch to the central processor. These minis could provide a small group (or office) with file-sharing, word-processing, spreadsheet, and e-mail, and manage output devices. They were marketed as “office systems,” “office automation (OA) systems,” or “office information systems” (OIS).

The 1980 Stanford International Symposium on Office Automation marked the emergence of a research field that remained influential for a decade and then faded away. Douglas Engelbart contributed two papers to the proceedings of this symposium (Landau, Bair, and Siegman 1982). In the same year, the American Federation of Information-Processing Societies (AFIPS, the parent organization of ACM and Institute of Electrical and Electronics Engineers [IEEE] at the time) held the first of seven annual OA conferences and product exhibitions. Also in 1980, ACM formed the Special Interest Group on Office Automation (SIGOA), which launched the biennial Conference on Office Information Systems (COIS) 2 years later. In 1983, the journal *ACM Transactions on Office Information Systems* (*TOOIS*) emerged, which was 1 year after the emergence of the independent journal *Office: Technology and People*.

You might ask “what is all this with offices?” Minicomputers brought down the price of computers to fit into the budget of a small workgroup or an office. (The attentive reader will anticipate: The personal computer era is approaching.) Office Information Systems, which focused on the use of minicomputers, was positioned alongside MIS, which focused on mainframes. Its scope was reflected in the charter of *TOOIS*: database theory, AI, behavioral studies, organizational theory, and communications. Minis were accessible to database researchers. Digital’s PDP series was a favorite with AI researchers until LISP machines flourished. Minis were familiar to behavioral researchers who used them to run and analyze psychology experiments. Computer-mediated communication (CMC) was an intriguing new capability: Networking was still rare, but people at different

terminals of a minicomputer could exchange e-mail or chat in real time. Minis became interactive computers of choice for many organizations. As a consequence, Digital became the second largest computer company in the world and Dr. Wang the fourth wealthiest American.

Researchers were discretionary users, but few office workers chose their tools. The term “automation” was challenging and exciting to researchers, but it conjured up less pleasant images for office workers. Some researchers, too, preferred Engelbart’s focus on augmentation rather than automation.

Papers in the SIGOA newsletter, COIS, and *TOOIS* included technical work on database theory, a modest number of AI papers (the AI winter had not yet ended), decision support and CMC papers from the IS community, and behavioral studies by researchers who later joined CHI. Papers on information systems were prevalent in the newsletter and technical papers in *TOOIS*, which also published numerous behavioral studies until the journal *Human-Computer Interaction* started in 1985.

Although OA/OIS research was eventually absorbed by other fields, it identified and called attention to important emerging topics, including hypertext, CMC, and collaboration support. OIS research was also allied with the technical side of information science, notably information retrieval and language processing.

THE FORMATION OF ASSOCIATION FOR COMPUTING MACHINERY SPECIAL INTEREST GROUP ON COMPUTER-HUMAN INTERACTION

Figure 1 identifies research fields that directly bear on HCI. Both HF and IS have distinct subgroups that focus on broad use of digital technologies. Relevant computer science research is concentrated in CHI, the subgroup primarily concerned with discretionary hands-on computer use. Other computer science influences—computer graphics, AI, office systems—have been described but are not included in Figure 1. The fourth field, information, began as support for specialists. It may come to exert the broadest influence of all.

Decreasing microcomputer prices encouraged discretionary hobbyists to use them. In 1980, as IBM prepared to launch the PC, a groundswell of attention on computer user behavior was building up. IBM, which like many hardware companies had not sold software separately, had decided to make software a product focus. Several cognitive psychologists joined an IBM group that included John Gould, who had been publishing human factors research since the late 1960s. They initiated empirical studies of programming and studies of software design and use. Other psychologists who in 1980 led recently formed HCI groups were Phil Barnard at the Medical Research Council Applied Psychology Unit in Cambridge, England; Tom Landauer at Bell Laboratories; Donald Norman at the University of California, San Diego; and John Whiteside at Digital Equipment Corp.

Xerox PARC and CMU collaborators continued research that led to an exceptionally influential project. The 1981 Star, with a carefully designed GUI, was not a commercial success




FIGURE 1 Four fields with major human–computer interaction research threads: Acronym expansions are provided in the text.

(nor were a flurry of GUIs that followed, including the Apple Lisa), but it influenced researchers and developers—and the design of the Macintosh.

Communications of the ACM created a “Human Aspects of Computing” department in 1980. The next year, Tom Moran edited a special issue of *Computing Surveys* on “The Psychology of the Computer User.” Also in 1981, the ACM Special Interest Group on Social and Behavioral Science Computing (SIGSOC) extended its workshop to cover interactive software design and use. In 1982, a conference in Gaithersburg, Maryland, on “Human Factors in Computing Systems” was unexpectedly well attended. Shortly afterward, SIGSOC shifted its focus to Computer-Human Interaction and changed its name to SIGCHI (Borman 1996).

In 1983, the first CHI conference attracted more than 1000 people. Half of the 58 papers were from the aforementioned seven research laboratories. Cognitive psychologists in industry dominated the program, although the Human Factors Society cosponsored the conference and contributed the program chair Richard Pew; committee members Sid Smith, H. Rudy Ramsay, and Paul Green; and several presenters. Brian Shackel and HFS president Robert Williges gave tutorials on the first day. The International Conference on Human–Computer Interaction (INTERACT), first held in London in 1984 and chaired by Shackel, drew HF&E and CHI researchers.

The first profession to become discretionary hands-on users was computer programming, as paper coding sheets were discarded in favor of text editing at interactive terminals, PCs, and small minicomputers. Therefore, many early

CHI papers, by Ruven Brooks, Bill Curtis, Thomas Green, Ben Shneiderman, and others, continued the psychology-of-programming research thread. Shneiderman formed the influential HCI Laboratory (HCIL) at Maryland in 1983. IBM researchers also contributed, as noted by John Thomas in a personal communication (October 2003): “One of the main themes of the early work was basically that we in IBM were afraid that the market for computing would be limited by the number of people who could program complex systems, so we wanted to find ways for ‘nonprogrammers’ to be able, essentially, to program.”

Many experimental psychologists undertook studies of text editing, a tool initially used primarily by programmers. Thomas Green remarked at INTERACT’84 that “text editors are the white rats of HCI.” As personal computing spread, studies of other discretionary use contexts were conducted. Studies of programming gradually disappeared from HCI conferences.

CHI focused on novice use. Initial experience is particularly important for discretionary users and for vendors developing software for them. Novice users are also a natural focus when studying new technologies and a critical focus when more people take up computing each year compared with the year before.

Routinized heavy use was still widespread. Databases were used by airlines, banks, government agencies, and other organizations. This hands-on activity was rarely discretionary. Managers oversaw development and analyzed data, leaving data entry and information retrieval to people hired for

those jobs. To improve data management tasks was a human factors undertaking. CHI studies of database use were few—I count three over a decade, all focused on novice or casual use.

Fewer European companies produced mass-market software. European HCI research focused on in-house development and use, as reflected in the journal *Behaviour & Information Technology*, which was launched in 1982 by Tom Stewart and published by Taylor & Francis in London. In his perceptive essay cited in the section “Discretion in Computer Use,” Bannon urged that more attention be paid to discretionary use, yet criticized CHI’s heavy emphasis on initial experience, reflecting the European perspective. At Loughborough University, HUSAT focused on job design (the division of labor between people and systems) and collaborated with the Institute for Consumer Ergonomics, particularly on product safety. In 1984, Loughborough initiated an HCI graduate program drawing on human factors, industrial engineering, and computer science.

The work of the early visionaries was unfamiliar to many CHI researchers who were helping realize some of the early visions. The 633 references in the 58 papers presented at CHI’83 included many authored by cognitive scientists, but Bush, Sutherland, and Engelbart were not cited. A few years later, more computer scientists familiar with the early work joined CHI, notably those working on interactive computer graphics. The psychologists eventually discovered and identified with the pioneers, who shared their concern for discretionary use. This conceptual continuity bestowed legitimacy on a young enterprise that sought to establish itself academically and professionally.

DIVERGENCE OF COMPUTER–HUMAN INTERACTION AND HUMAN FACTORS

Hard science, in the form of engineering, drives out soft science, in the form of human factors.

Newell and Card (1985, p. 212)

Between 1980 and 1985, Card, Moran, and Newell (1980a,b) introduced a “keystroke-level model for user performance time with interactive systems,” followed by the cognitive model goals, operators, methods, and selection rules (GOMS) in their landmark 1983 book *The Psychology of Human–Computer Interaction*. This work was highly respected by the cognitive psychologists prevalent in CHI at the time. However, these models did not address discretionary, novice use. They focused on the repetitive expert use studied in human factors. In fact, GOMS was explicitly positioned to counter the latter field’s stimulus–response bias: “Human-factors specialists, ergonomists, and human engineers will find that we have synthesized ideas from modern cognitive psychology and AI with the old methods of task analysis. . . . The user is not an operator. He does not operate the computer, he communicates with it” (Newell and Card 1985, p. viii.).

Newell and Card noted that HF’s had a role in design, but continued: “Classical human factors . . . has all the earmarks of second-class status. (Our approach) avoids continuation of the classical human-factors role (by transforming) the psychology of the interface into a hard science” (p. 221).

In 2004, Card noted in an e-mail discussion: “Human Factors was the discipline we were trying to improve. . . . I personally changed the (CHI conference) call in 1986, so as to emphasize computer science and reduce the emphasis on cognitive science, because I was afraid that it would just become human factors again.”

Ultimately, human performance modeling drew a modest but fervent CHI following. Key goals differed from those of other researchers and many practitioners. “The central idea behind the model is that the time for an expert to do a task on an interactive system is determined by the time it takes to do the keystrokes,” wrote Card, Moran, and Newell (1980b, p. 397). Modeling was extended to a range of cognitive processes, but it was most useful in helping to design for non-discretionary users such as telephone operators engaged in repetitive tasks (e.g., Gray et al. 1990). Its role in augmenting human intellect was unclear.

CHI and HFS moved apart, although “Human Factors in Computing Systems” remains the CHI conference subtitle. They were never highly integrated. Most of the cognitive psychologists had turned to HCI after earning their degrees and were unfamiliar with the human factors literature. The Human Factors Society did not again cosponsor CHI. Its researchers disappeared from the CHI program committee. Most CHI researchers who previously published in the human factors literature shifted to CHI, *Communications of the ACM*, and the journal *Human–Computer Interaction* launched in 1985 by Thomas Moran and published by Erlbaum, a publisher of psychology books and journals.

The shift was reflected at IBM T.J. Watson Research Center. John Gould and Clayton Lewis authored a CHI’83 paper that nicely framed the CHI focus on user-centered, iterative design based on prototyping. Cognitive scientists at Watson helped shape CHI, but Gould’s principal focus remained human factors; he served as HFS president 4 years later. Reflecting the broader change, in 1984 the Human Factors Group at Watson began to dissolve and a User Interface Institute emerged.

CHI researchers, identifying with “hard” science or engineering, adopted the terms “cognitive engineering” and “usability engineering.” In the first paper presented at CHI’83, “Design Principles for Human–Computer Interfaces,” Donald Norman (1983) applied engineering techniques to discretionary use, creating “user satisfaction functions” based on technical parameters. These functions would not hold up long—people are fickle, yesterday’s satisfying technology is not as gratifying today—but for years CHI emulated engineering, downplaying design, marketing, and other aspects of how humans interact with technology.

WORKSTATIONS AND ANOTHER ARTIFICIAL INTELLIGENCE SUMMER

High-end workstations from Apollo, Sun, and Silicon Graphics appeared between 1981 and 1984. Graphics researchers no longer had to flock to heavily financed laboratories (notably MIT and Utah in the 1960s; MIT, New

York Institute of Technology, and PARC in the 1970s). Workstations were too expensive to reach a mass market, so graphics research that focused on photorealism and animation, which required the processing power of workstations, did not directly exert a broad influence on HCI.

The Xerox Star (formally named Office Workstation), Apple Lisa, and other commercial GUIs appeared, but when the first CHI conference was held in December 1983 none were commercial successes. They cost too much or ran on processors that were too weak to exploit graphics effectively.

In 1981, Symbolics and LMI introduced workstations optimized for the LISP programming language favored by most AI researchers. The timing was fortuitous. In October of that year, a conference on next-generation technology was held in the National Chamber of Commerce auditorium in Tokyo, Japan, and in 1982 the Japanese government established the Institute for New Generation Computer Technology (ICOT) and a 10-year fifth generation project focused on AI. AI researchers in Europe and the United States sounded the alarm. Donald Michie of Edinburgh saw a threat to Western computer technology, and in 1983 Ed Feigenbaum of Stanford and Pamela McCorduck wrote: “The Japanese are planning the miracle product.... They’re going to give the world the next generation—the Fifth Generation—of computers, and those machines are going to be intelligent.... We stand, however, before a singularity, an event so unprecedented that predictions are almost silly.... Who can say how universal access to machine-intelligence—faster, deeper, better than human intelligence—will change science, economics, and warfare, and the whole intellectual and sociological development of mankind?” (pp. 8–9, 287).

Parallel distributed processing (often called neural networks) models also seized the attention of researchers and the media. Used for modeling phenomena including signal detection, motor control, and semantic processing, neural networks represented conceptual and technical advances over earlier AI work on perceptrons. Their rise was tied to the new generation of minicomputers and workstations, which had the power to support simulation experiments. Production systems, a computer-intensive AI modeling approach with a psychological foundation, developed at CMU, also gained the attention of researchers.

These developments triggered an AI gold rush. As with actual gold rushes, most of the money was made by those who outfitted and provisioned the prospectors, although generous government funding again flowed to the actual researchers. The European ESPRIT and UK Alvey programs invested over US\$200 million per year starting in 1984 (Oakley 1990). In the United States, funding for the DARPA Strategic Computing AI program, begun in 1983, rose to almost US\$400 million in 1988 (Norberg and O’Neill 1996). Investment in AI by 150 U.S. companies was estimated at about US\$2 billion in 1985 (Kao 1998).

The unfulfilled promises of the past led to changes this time around. General problem solving was emphasized less, whereas domain-specific problem solving was emphasized

more. Terms such as intelligent knowledge-based systems, knowledge engineering, expert systems, machine learning, language understanding, image understanding, neural networks, and robotics were often favored over AI.

In 1983, Raj Reddy of CMU and Victor Zue of MIT criticized the mid-1970s abandonment of speech-processing research, and soon funds again became plentiful for these research topics (Norberg and O’Neill 1996, p. 238). Johnson (1985) estimated that 800 corporate employees and 400 academics were working on natural language-processing research. Commercial natural language-understanding (NLU) interfaces to databases such as AI Corporation’s Intellect and Microrim Clout appeared.

The optimism is illustrated by two meticulously researched Ovum reports on speech and language processing (Johnson 1985; Engelen and McBride 1991). In 1985, speech and language product “revenue” was US\$75 million, comprising mostly income from grants and investor capital. That year, Ovum projected that sales would reach US\$750 million by 1990 and US\$2.75 billion by 1995. In 1991 sales were under US\$90 million, but hope springs eternal and Ovum forecasts US\$490 million for 1995 and US\$3.6 billion for 2000.

About 20 U.S. corporations banded together, jointly funding the Microelectronics and Computer Technology Corporation (MCC). U.S. antitrust laws were relaxed to facilitate this cooperation. MCC embraced AI, reportedly becoming the leading customer for both Symbolics and LMI. MCC projects included two parallel NLU efforts; work on intelligent advising; and CYC (as in encyclopedic, and later spelled Cyc), Douglas Lenat’s ambitious project to build a commonsense knowledge base that other programs could exploit. In 1984, Lenat predicted that by 1994 CYC would be intelligent enough to educate itself. Five years later, CYC was reported to be on schedule and about to “spark a vastly greater renaissance in [machine learning]” (Lenat 1989, p. 257).

Knowledge engineering involved human interaction. This could have brought AI closer to HCI, but AI researchers who were interested in representation and reasoning were frustrated by the difficulty of eliciting knowledge from experts. As many AI systems were aimed at nondiscretionary use, this created opportunities for HF&E, especially in Europe where funding directives dictated work that spanned technical and behavioral concerns. The journal *IJMMIS* became a major outlet for both HF&E and AI researchers in the 1980s.

Interaction of AI and CHI was limited. CHI’83 and CHI’85 had a few sessions on speech and language, cognitive modeling, knowledge-based help, and knowledge elicitation. Not many AI researchers and developers worried about usability. They loved powerful tools such as EMACS and UNIX, forgetting the painful weeks required to learn the badly designed command languages. In general, AI technologies did not succeed in the marketplace. Before it disappeared, AI Corporation’s primary customer for the database interface Intellect was the government, where discretionary use was not the norm.

1985–1995: GRAPHICAL USER INTERFACES SUCCEED

“There will never be a mouse at the Ford Motor Company.”
A high-level acquisition manager, 1985
 (personal communication)

When graphical user interfaces finally succeeded commercially, human–computer interaction was transformed. As with previous disruptive shifts—to stored programs and to interaction based on commands, full-screen forms, and full-screen menus—some people were affected before others. GUIs were particularly attractive to consumers, to new or casual users. Their success immediately transformed CHI, but only after Windows 3.0 succeeded in 1990 did GUIs influence the government agencies and business organizations that are the focus of HF&E and IS researchers. By 1990, the technology was better understood and thus less disruptive. The early 1990s also saw the maturation of local area networking and the Internet, producing a second transformation: computer-mediated communication and information sharing.

COMPUTER–HUMAN INTERFACE EMBRACES COMPUTER SCIENCE

Apple launched the Macintosh with a 1984 Super Bowl ad describing office work, but sales did not follow and by mid-1985 Apple was in trouble. Then Macs appeared with four times as much random access memory (RAM), which was sufficient to manage Aldus PageMaker, Adobe Postscript, the Apple LaserWriter, and Microsoft’s Excel and Word for Macintosh as they were released. The more powerful Mac Plus arrived in January 1986. Rescued by hardware and software advances, the Mac succeeded where many commercial GUIs before it could not. It was popular with consumers and became the platform for desktop publishing.

Within CHI, GUIs were initially controversial. They had disadvantages: An extra level of interface code increased development complexity and created reliability challenges. They consumed processor cycles and distanced users from the underlying system that, many believed, experienced users must eventually master. Carroll and Mazur (1986) showed that GUIs confused and created problems for people familiar with existing interfaces. An influential 1986 essay on direct manipulation interfaces by Hutchins, Hollan, and Norman concluded that “it is too early to tell” how GUIs would fare. The GUIs could well prove useful for novices, they wrote, but “we would not be surprised if experts are *slower* with Direct Manipulation systems than with command language systems” (pp. 119-121, italics in the original). Given that most prior HCI research had focused on expert use, this insight seemed significant. However, first-time use proved critical in the rapidly expanding consumer market, and hardware and software improvements overcame some early limitations. GUIs were here to stay. CHI was soon transformed. Previously active research topics, including command naming, text editing, and the psychology of programming, were

abandoned. More technical topics such as “user interface management systems” became significant.

Viewed from a higher plane, psychology gave way to computer science as the driving force in interaction design. Researchers had strived for a comprehensive, theoretical, psychological framework based on formal experiments (Newell and Card 1985; Carroll and Campbell 1986; Long 1989; Barnard 1991). Such a framework was conceivable for constrained command- and form-based interaction but could not be scaled to design spaces that included color; sound; animation; and an endless variety of icons, menu designs, and window arrangements. The new mission was to identify the most pressing problems and find satisfactory rather than optimal solutions. Rigorous experimentation, a skill of cognitive psychologists, gave way to quicker, less precise assessment methods championed by Jakob Nielsen (1989; Nielsen and Molich 1990).

Exploration of the dynamically evolving, relatively unconstrained design space required software engineering expertise. The late 1980s saw an influx of computer scientists to the CHI community. HCI entered the curricula of many computer science programs. CHI became a natural home to some computer scientists working on interactive graphics, software engineers interested in interaction, and AI researchers working on speech recognition, language understanding, and expert systems. In 1994, ACM launched the journal *Transactions on Computer–Human Interaction (TOCHI)*. Early PCs and Macs were not easily networked, but as the use of local area networks spread, CHI’s focus expanded to include collaboration support. This brought it into contact with efforts in MIS and OA research, discussed in the section on Collaboration Support below.

HUMAN FACTORS AND ERGONOMICS MAINTAINS A NONDISCRETIONARY USE FOCUS

Human factors and ergonomics research continued to respond to the needs of government agencies, the military, aviation industry, and telecommunications. Governments are the largest consumers of computing, for census, tax, social security, health and welfare, power plant operation, air traffic control, ground control for space missions, military logistics, text and voice processing for intelligence, and so on. The focus is on skilled use—users are assigned technology and trained if necessary. For routine data entry and other tasks, small efficiency gains in individual transactions can yield large benefits over time, justifying the effort to make improvements that might not be noticed by discretionary users. After SIGCHI formed, HFS undertook a study to see how CHI would affect membership in its Computer Systems Technical Group. An unexpectedly small effect was found (Richard Pew, personal communication; September 15, 2004). They had different goals.

Government agencies promoted the development of ergonomic standards to help in defining system requirements for competitive bidding while remaining at arms’ length from potential developers, who of course better understood technical possibilities and helped with standards development.

Compliance with standards could then be specified in a contract. In 1986, Sid Smith and Jane Mosier published the last of a series of government-sponsored interface guidelines, with 944 design guidelines organized into sections titled Data Entry, Data Display, Data Transmission, Data Protection, Sequence Control, and User Guidance. The authors recognized that GUIs would expand the design space beyond the reach of this already cumbersome document that omitted icons, pull-down and pop-up menus, mice button assignments, sound, animation, and so on. Smith and Mosier foresaw that requirements definition must shift to specify predefined interface styles and design processes rather than features that would be built from scratch.

DARPA's heavily funded strategic computing AI program set out to develop an autonomous land vehicle, a pilot's associate, and a battle management system. All raised human factors research issues. These systems were to include interactive technologies such as speech recognition, language understanding, and heads-up displays. People might avoid these technologies when given a choice, but pilots guiding autonomous vehicles and officers under stressful conditions might have no better alternative. Speech and language technologies have other nondiscretionary potential, some of it civilian: for translators and intelligence analysts, when a phone system provides no alternative, when a disability limits keyboard use, or when hands are otherwise occupied.

INFORMATION SYSTEMS EXTENDS ITS RANGE

Although GUIs were not quickly adopted by organizations, spreadsheets and business graphics (charts and tables) were important to managers and thus the foci of IS research. Remus (1984) contrasted tabular and graphic presentations and Benbasat and Dexter (1985) added color as a factor, although color displays were rare in the 1980s. Many studies contrasted online and paper presentations, because most managers worked with printed reports. Although research into individual cognitive styles was abandoned in the early 1980s following a devastating critique on the topic (Huber 1983), the concept of cognitive fit between task and tool was introduced to explain apparently contradictory results in the adoption literature (Vessey and Galletta 1991).

A series of symposia on human factors in IS was initiated in 1986 by Jane Carey, leading to several books on the subject (e.g., Carey 1988). Topics included user interaction with information, design and development and, as corporate adoption of minicomputers and intranets matured, communication and collaboration, including studies of e-mail use.

The involvement of end users in the development process was actively discussed in IS, but rarely practiced outside of the sociotechnical design and the participatory design movements discussed below in the section "Participatory Design and Ethnography" (Friedman 1989). Hands-on managerial use was atypical in this period, but it was central to group decision support systems (GDSS) research. Central to GDSS was support for meetings, including brainstorming, idea organization, and online voting features. GDSS emerged from

decision support systems, aimed at supporting individual executives or managers, and later evolved into group support systems. Computer-supported meeting facility research was conducted in the mid-1980s in several laboratories (e.g., Begeman et al. 1986; DeSanctis and Gallupe 1987; Dennis et al. 1988). Extensive research at the University of Arizona is summarized by Nunamaker et al. (1997). These systems were initially too expensive to be mass-market products; hence, the focus was on "decision makers," and research was conducted primarily in schools of management, not computer science departments or software companies. GDSS was a major IS contribution to computer-supported cooperative work (CSCW), discussed in the next section. In 1990, three companies began marketing GDSSs, including IBM and a University of Arizona spin-off, although without much success.

The Technology Acceptance Model (TAM) introduced by Davis (1989) led to considerable IS research. TAM and its offspring focus on perceived usefulness and perceived usability to improve "white-collar performance" that is "often obstructed by users' unwillingness to accept and use available systems" (p. 319). "An element of uncertainty exists in the minds of decision makers with respect to the successful adoption," wrote Bagozzi, Davis, and Warshaw (1992, p. 664). Although TAM is a managerial view of individual behavior, it was influenced by Davis's exposure to early CHI usability research.

TAM is probably the most cited HCI work in IS. The management view of hands-on computer use as nondiscretionary was giving way as use spread to white-collar workers who could refuse to play. TAM's emphasis on *perceived* utility and usability is a key distinction: Consumers choose technologies that they are convinced will be useful; CHI researchers assume utility and focuses on the *experience* of usability. TAM researchers focus on utility and note that perceptions of usability can influence acceptance. CHI addressed usability a decade before TAM, albeit actual usability rather than perceived usability. Perception was a secondary 'user satisfaction' measure to CHI researchers, who believed (not entirely correctly) that measurable reduction in time, errors, questions, and training would eventually translate into positive perceptions. The word "acceptance," that is, the "A" in TAM, is not in the CHI vocabulary. Discretionary users adopt, they do not accept.

The IS and CHI communities rarely mixed. When CHI was over a decade old, *Harvard Business Review*, a touchstone for IS researchers, published "Usability: The New Dimension of Product Design" (March 1994). The article did not mention CHI at all. It concluded that "user-centered design is still in its infancy" (p. 149).

COLLABORATION SUPPORT: OFFICE INFORMATION SYSTEMS GIVES WAY TO COMPUTER-SUPPORTED COOPERATIVE WORK

In the late 1980s, three research communities addressed small-group communication and information sharing: (1) OA/OIS, described above in the section "Minicomputers

and Office Automation.” (2) IS researchers building systems to support organizational decision making could, as computing costs declined, address group decision making more generally. (3) The proliferation of local area networks enabled some CHI researchers to move from individual productivity software to the quest for “killer apps” that would support teams.

OA/OIS led the way, but it declined and was fast disappearing by 1995. The Minicomputers, the platform for most OIS research, did not survive competition from PCs and workstations. The concept of office or group proved to be problematic: Organizations and individuals are persistent entities with goals and needs, but small groups often have ambiguous membership and undergo shifts in character as members join or depart. People in an organization who need to communicate often fall under different budgets, complicating acquisition decisions unless a technology is made available organization-wide.

The rapid shift was reflected in terminology use. First, “automation” fell out of favor. In 1986, ACM SIGOA shifted to SIGOIS and the annual AFIPS OA conferences were discontinued. By 1991, the term “office” followed: *Transactions on Office Information Systems* became *Transactions on Information Systems*; *Office: Information and People* became *Information Technology and People*; and “Conference on Office Information Systems” became “Conference on Organizational Communication Systems” (COOCS, in 1997 becoming the GROUP Conference).

The AI summer, which contributed to the OA/OIS effort, ended when AI failed to meet expectations: Massive funding did not deliver a pilot’s associate, an autonomous land vehicle, or a battle management system for the military. Nor were offices automated. CHI conference sessions on language processing had diminished prior to this AI winter, but sessions on modeling, adaptive interfaces, advising systems, and other uses of intelligence in interfaces increased through the 1980s before declining in the 1990s. Funding for AI became scarce, employment opportunities dried up, and conference participation dropped off.

A 1986 conference, building on a successful private 1984 workshop (Greif 1985), brought together researchers from diverse disciplines interested in issues of communication, information sharing, and coordination under the banner “Computer Supported Cooperative Work.” Participants came primarily from IS, OIS, CHI, distributed AI, and anthropology. Four of 13 CSCW program committee members and many papers were from schools of management, with similar participation by the OIS community.

The field coalesced in 1988. The book *Computer-Supported Cooperative Work*, edited by Irene Greif, was published, and SIGCHI sponsored a biennial North American CSCW conference. A European series (ECSCW) was initiated in 1989. With heavy participation from technology companies, North American CSCW had a small-group focus on networked individuals working on PCs, workstations, or minicomputers. Groups were either within an organization or linked by ARPANET, BITNET, or other networks. European

participation, primarily from academia and government agencies, focused on organizational use of technologies. It differed methodologically from most IS research in North America. Scandinavian influences, described in the next section, were felt in both CSCW and ECSCW.

Just as human factors researchers left CHI after a few years, most IS researchers who were involved with CSCW left in the early 1990s. One factor was a shift within IS from social psychology to organizational behavior in studying team behavior. The Hawaii International Conference on System Sciences (HICSS) was becoming a major IS prejournal publication venue for work with an organizational orientation. In contrast, the organizational focus conflicted with the CSCW interest in context-independent small-group support, which was the realm of social psychology and the goal of many technology companies. Some IS researchers participated in COOCS and GROUP. The split was not entirely amicable; the IS newsletter *Groupware Report* did not include CSCW on its list of relevant conferences.

The pace of technology change created challenges for CSCW. In 1985, supporting a small team was a technical challenge; 10 years later, the Web had arrived. Applications that provided awareness of the activity of distant collaborators was a celebrated achievement in the early 1990s; several years later, dark linings to the silver cloud arose in the form of privacy concerns and information overload. Phenomena, such as a “productivity paradox” in which IT investments were not returning benefits and health effects of Internet use by young people, were carefully identified only to vanish a few years later. Other changes brought European and North American CSCW into greater alignment. European organizations were starting to acquire commercial software products, a CSCW focus in North America, and North Americans were discovering that organizational context, an ECSCW focus, was often crucial in the design and deployment of products intending to support group activity. Organizational behaviorists and theorists were thriving in their home disciplines, but ethnographers studying technology use, marginalized in traditional anthropology departments, were welcomed into CSCW.

Despite the challenges of building on sands swept by successive waves of technology innovation, CSCW remains a strong research area that attracts a broad swath of HCI researchers. Content ranges from the highly technical to thick ethnographies of workplace activity, from studies of instant messaging dyads to scientific laboratories involving hundreds of people dispersed in space and time. Chapter 24 by Gary and Judy Olson in this handbook covers the technical side of this topic in depth, with references to other CSCW resources.

PARTICIPATORY DESIGN AND ETHNOGRAPHY

Prior to 1985-1995 some system developers explored methods to involve some of the future users in designing a system. Typically the users were nondiscretionary users of a system being developed by a large enterprise for its own use.

Sociotechnical design took a managerial perspective. Participatory or cooperative design, rooted in the Danish trade union movement, focused on empowering eventual users (Nygaard 1977).

Scandinavian approaches influenced human factors (e.g., Rasmussen 1986) and attracted wide notice with the publication of the proceedings of a conference held in Aarhus, Denmark, in 1985 (Bjerknes et al. 1987). Participatory design was a critique of IS approaches, yet the Scandinavians resonated with CHI researchers. Despite differences in culture, contexts of development (in-house system vs. commercial product), and contexts of use (nondiscretionary vs. discretionary), they shared the goal of empowering hands-on users. Most were also of the generation that grew up in the 1960s, unlike the World War II generation that dominated HF&E and IS.

Ethnography was a different approach to obtaining deep insights into potential users. Lucy Suchman managed a Xerox PARC group that presented studies of workplace activity at CSCW. Suchman published an influential critique of artificial intelligence in 1987 and a widely read review of the Aarhus proceedings in 1988, and as program chair she brought many Scandinavians to the CSCW 1988 conference.

LIBRARY AND INFORMATION SCIENCE: AN INCOMPLETE TRANSFORMATION

Research universities have always supported prestigious professional schools, but the prestige of library schools declined with the rise of higher-paid IT and software engineering professions. Between 1978 and 1995, 15 American library schools were shut down (Cronin 1995, p. 45). Most of the survivors were rechristened Library and Information Science. The humanities orientation had given way, and librarianship was being changed by technology. New curricula and faculty with different skills were needed.

The changes did not go smoothly or as anticipated. Forced multidisciplinary is never easy. Exclusion of technology studies may have been a reasonable reaction to the expense and limitations of new technologies. However, Moore's law lowered costs and removed many limitations with such speed that people and organizations had little time to prepare. Young information scientists were not interested in absorbing a century of work on indexing, classifying, and providing access to complex information repositories; their eyes were fixed on a future in which many past lessons would not apply. Those that still applied would likely have to be relearned. The conflicts are exposed in a landmark 1983 collection, *The Study of Information: Interdisciplinary Messages* (Machlup and Mansfield 1983). In the book, W. Boyd Rayward outlines the humanities-oriented perspective and the technological perspective and argues that there was convergence. His essay is followed by commentaries attacking him from both sides.

In a series of meetings beginning in 1988, new library and information school deans at the universities Pittsburgh, Syracuse, Drexel, and subsequently Rutgers discussed approaches to explaining and managing multidisciplinary

schools. Despite this progressive effort, Cronin (1995) depicted LIS at loggerheads and in a "deep professional malaise." He suggested that librarianship be cut loose in favor of stronger ties to cognitive and computer sciences. Through the 1990s, schools at several universities dropped the word "library" and became schools of information (see Figure 2). More would follow.

1995–2010: THE INTERNET ERA ARRIVES

How did the spread of the Internet and the emergence of the Web affect HCI research threads? CHI researchers were Internet savvy. Although excited by the prospects, they took these changes in stride. Over time, CHI-related research, development, and use evolved. The Internet and the Web were not disruptive to HF&E either. The Web was initially a return to a form-driven interface style, and it was rarely a locus of routine work. However, the Web had a seismic impact on IS and on information science, so this section begins with these disciplines.

THE FORMATION OF ASSOCIATION FOR INFORMATION SYSTEMS SPECIAL INTEREST GROUP IN HUMAN–COMPUTER INTERACTION

The use of computers in organizations has changed. Organizations are no longer focused on maximizing computer use—almost everywhere, screen savers have become the main consumer of processor cycles. Advent of the Internet created more porous organizational boundaries. Employees in many organizations could download software such as instant-messaging clients, music players, and weblog tools inside organizational firewalls despite IT concerns about productivity and security. These are not the high-overhead applications of the past. Increasingly, software can be used from a web browser without requiring a download. Experience with all of this at home leaves employees impatient with poor software at work. In addition, many managers who had been hands-off users became late adopters in late 1990s or were replaced by younger managers. Today, managers and executives are hands-on early adopters of many technologies.

Significant as these changes are, the Web had a more dramatic effect on organizational information systems. Corporate IT groups had been focused solely on internal operations. They lived inside firewalls. Their customers were other employees. Suddenly, organizations were scrambling to create Web interfaces to external vendors and customers. Discretionary users! The Internet bubble burst, revealing that IT professionals, IS experts, and everyone else had limited understanding of Web phenomena. Nevertheless, online marketing, services, and business-to-business systems continued to grow. For many, the Web had become an essential business tool. In handling external customers, IT professionals and IS researchers were in much the same place that CHI was 20 years earlier, whether they realized it or (most often) not.

In 2001, the Association for Information Systems (AIS) established a Special Interest Group in Human–Computer




FIGURE 2 The iSchools and when "information" came into the names of the member Schools (Faculties, Colleges, etc.).

Interaction (SIGHCI). The founders defined HCI by citing 12 CHI research papers (Zhang, Nah, and Preece 2004, p. 148). Bridging the CHI and the information science communities was declared a priority. The charter of SIGHCI includes a broad range of organizational issues, but the publications emphasize interface design for e-commerce, online shopping, online behavior "especially in the Internet era" (Zhang 2004, p. 1), and effects of Web-based interfaces on attitudes and perceptions. Eight of the first 10 papers in SIGHCI-sponsored journal issues covered Internet and Web behavior.

In 2009, the journal *AIS Transactions on Human-Computer Interaction* was launched. The shift from an organizational focus to the Web and broader end-user computing is documented in Zhang et al.'s analysis (2009) of the IS literature from 1990 to 2008. This survey omits CHI from a list of the fields related to AIS SIGHCI. The bridging effort had foundered, as had three previous efforts to bridge to CHI: from Human Factors, Office Information Systems, and the Information Systems presence within CSCW.

DIGITAL LIBRARIES AND THE EVOLUTION OF LIBRARY INFORMATION SCIENCE

By 1995, an information wave had swept through universities (Figure 2). Digital technology was in the LIS curriculum. Familiarity with technology use was a prerequisite for librarianship. However, innovative research had not kept pace with professional training (Cronin 1995).

The Internet grew exponentially, but in 1995 it was still a niche activity found mainly on campuses. In the mid-1990s, Gopher, a convenient system for downloading files over the Internet, attracted attention as a possible springboard for indexing distributed materials. Wells's (1938) concept of "world brain" seemed to be within reach. Then the Web hit, transforming information acquisition, management, and access at an ever-increasing pace. Between 1994 and 1999, two NSF/DARPA/NASA/National Library of Medicine/Library of Congress/National Endowment for the Humanities/FBI initiatives awarded close to US\$200 million for digital libraries research and development. This and other

investments galvanized the research community. In 2000, the American Society for Information Science appended “and Technology” to its name to become ASIST.

By 2000, 10 schools (or equivalent units) had information as the sole discipline in their name. In 2001 a series of deans meetings began, which were modeled on those of the late 1980s. The original members, Syracuse, Pittsburgh, and Drexel, were joined by Michigan; Berkeley, California; and the University of Washington. All are now information schools. In 2005, the first annual “iConference” drew participants from 19 universities with information programs. As of 2011, the “iCaucus” had 27 dues-paying members. Some are transformed library schools, some have closer ties with other disciplines, and some have formed recently as schools of information. Collectively, their faculty includes HCI researchers trained in each of the four disciplines highlighted in this introduction.

Expansion is not without growing pains. Conflicts arise among academic subcultures. The iConference competes with more established conferences in each field. Figure 2 suggests that a shift to a field called information is well underway, but many faculty still consider themselves “a researcher in {X} who is located in an information school,” where X could be library science, HCI, CSCW, IS, communication, education, computer science, or another discipline. We do not know how it will evolve, but we can say with confidence that information has become, and will remain, a significant player in HCI.

HUMAN FACTORS AND ERGONOMICS EMBRACES COGNITIVE APPROACHES

In 1996, the HFES formed a new technical group, Cognitive Engineering and Decision Making. It quickly became the largest technical group. A decade earlier this would have been unthinkable: Some leading human factors researchers disliked cognitive approaches. The CHI community first used the term cognitive engineering in this sense (Norman 1982, 1986). As this development suggests, CSTG declined in size and prominence as the HCI community dispersed. Most HF&E technical groups, from groups on telecommunications to those on medical systems, address digital technology and thereby HCI-related research.

Equally astonishing, in 2005 Human Performance Modeling was a new and thriving HFES technical group, initiated by Wayne Gray and Dick Pew, who had been active in CHI in the 1980s. Card, Moran, and Newell (1983) had introduced human performance modeling to reform the discipline of Human Factors from without. Some work continued within CHI that was focused on expert performance (e.g., a special issue of *Human–Computer Interaction*, vol. 12, number 4, 1997), but today the reform effort has moved within HF&E and remains focused largely on nondiscretionary use.

Government funding of HCI was largely shaped by the focus of HF&E. The Interactive Systems Program of the U.S. NSF—subsequently renamed HCI—was described thus: “The Interactive Systems Program considers scientific

and engineering research oriented toward the enhancement of human–computer communications and interactions in all modalities. These modalities include speech/language, sound, images and, in general, any single or multiple, sequential, or concurrent, human–computer input, output, or action” (National Science Foundation 1993).

One NSF program manager identified his proudest accomplishment to be doubling the already ample funding for natural language understanding research. Even after NSF established a separate Human Language and Communication Program in 2003, speech and language research was heavily supported by both the HCI and accessibility programs, with lighter support from AI and other programs. Subsequent NSF HCI program managers emphasized “direct brain interfaces” or “brain–computer interaction” based on brain waves and implants. A review committee noted that a random sample of NSF HCI grants included none by prominent CHI researchers (National Science Foundation 2003). NSF program managers rarely attended CHI conferences, which have little coverage of speech, language, or direct brain interaction. These technologies may prove useful, but they have so far made few inroads into discretionary use situations in homes and offices.

COMPUTER–HUMAN INTERACTION EVOLVES, AND EMBRACES DESIGN

The steady flow of new hardware, software features, applications, and systems ensures that people are always encountering and adopting digital technologies for the first time. This is important for technology producers and it generates new research issues. CHI has tracked this, generally focusing on an innovation when it first starts to attract a wide audience.

As an application matures, its use often becomes routine. Technologies such as e-mail and word processing, no longer discretionary for most of us, get less attention from CHI researchers whose gaze is directed toward the discretionary use of the moment, including Web design, ubiquitous and mobile computing, social computing, and use of Wikipedia. New issues include information overload, privacy, and effects of multitasking, and encourage the emergence of new methods, such as ethnography and data mining. At a higher level, continuity is found in CHI: exploration of input devices, communication channels, information visualization techniques, and design methods. Proposals to build HCI theory on these shifting sands (Barnard et al. 2000; Carroll 2003) remain largely aspirational.

Expanding participation in the Internet as its reliability and bandwidth increased steadily through the mid-1990s brought real-time and quasi-real-time communication technologies such as e-mail into greater focus. The Web temporarily slowed this by shifting attention to indirect interaction with static sites, but with the advent of Web 2.0 and greater support for animation and video the pace quickened. The Web was like a new continent. Explorers posted flags here and there. Then came attempts at settlement, with the virtual worlds research and development that blossomed in the late

1990s. Few of the early pioneers survived; there was little to do in virtual worlds other than chat and play games. But slowly some people shifted major portions of their work and play online, relying on online information sources, digital photo management, social software, digital documents, online shopping, multiplayer games, and so on. This evolution is reflected in CHI research.

The content of CSCW in North America has shifted in response to the extraordinary growth of social networking sites, Wikipedia, and other Web phenomena, which are of intense interest to students and academic researchers and the software companies who hire or consult with many of them. These technologies are not yet of great interest to the organizations and government agencies that are the customer for European CSCW research, and the move toward shared interests has been reversed. Europeans have moved more rapidly into basic research in vertical domains. The division resembles that of 20 years ago, based on a new generation of technology. In several years the two research threads may again converge, perhaps under different names: “computer supported cooperative work” is outdated. Many digital devices are not considered computers, they play central rather than support roles, activities around them can be competitive or conflictual, and they may be used more for recreation than work.

The Web curtailed research into one thread of AI research: powerful, self-contained personal productivity tools. Considerable effort is required to embed knowledge in application software, but when access to external information sources was limited, it was worth trying. With today’s easy access to information and knowledgeable people online, static, self-contained knowledge representation is less useful. In contrast, adaptive systems that merge and filter local and Internet-based information have a role to play. Steady progress in machine learning is enhancing productivity tools, although implausible AI forecasts have not disappeared.

To the psychologists and computer scientists who formed the CHI community, interface design was a matter of science and engineering. They focused on performance and assumed that people eventually choose efficient alternatives. Because human discretion involves aesthetic preferences and invites marketing and nonrational persuasion, this view was not sustained when computing costs came down. This engineering orientation gripped CHI longer than SIGGRAPH, where aesthetic appeal was a major driver. CHI researchers eventually came around, labeling the study of enjoyment “funology” (Blythe et al. 2003) lest someone think that they were having too good a time.

Some visual designers participated in graphical interface research early on. Aaron Marcus began working full time on computer graphics in the late 1960s. William Bowman’s book *Graphic Communication* (1968) was a strong influence on the development of Xerox Star, for which the designer Norm Cox’s icons were chosen (Bewley et al. 1983). However, graphic design was considered a secondary activity (Evenson 2005). In 1995, building on workshops at previous conferences, SIGCHI initiated “Designing Interactive Systems”

(DIS), a biennial conference that draws more systems designers than visual designers. In 2003, SIGCHI, SIGGRAPH, and the American Institute of Graphic Arts (AIGA) initiated the “Designing for User Experience” (DUX) conference series that fully embraced visual and commercial design. This effort lasted only through 2007, but the significance of design was established. Design is not typically assessed in research papers. The changing sensibility is reflected in *ACM Interactions*, a magazine launched by CHI in 1994, which has steadily increased the focus on design in both its content and its appearance.

Design’s first cousin, marketing, has been poorly regarded by the CHI community (Marcus 2004). Website design forced the issue. Site owners wish to keep users interested in a site, whereas users may prefer to escape quickly. Consider supermarkets, which position items that most shoppers want far apart, forcing people to traverse aisles where other products beckon. CHI professionals who align themselves with end users face a stakeholder conflict when designing for a site owner. This was not true in the past: Designers of individual productivity tools had little conflict of interest with prospective customers. Marketing is concerned with identifying and satisfying user needs, as well as shaping them. It will likely find a place in CHI, perhaps labeled “brandology.”

Finally, CHI has gradually become more open to work that takes a social or political stance. Accessibility was first addressed in the context of physical constraints. Socioeconomic factors were included in Universal Usability conferences in 2000 and 2003. Sustainability and fitness emerged as topics. This may reflect a distancing from a sense that engineering should strive for value neutrality, a bid for relevance by an increasingly academic group or aging CHI baby boomers who are considering their legacies.

The evolution of CHI is reflected in the influential contributions of Donald Norman. A cognitive scientist who introduced the term cognitive engineering, he presented the first CHI’83 paper. It defined “user satisfaction functions” based on speed of use, ease of learning, required knowledge, and errors. His influential book *Psychology of Everyday Things* (1988) focused on pragmatic usability. Its 1990 reissue as *Design of Everyday Things* reflected a field refocusing on invention. Fourteen years later he published *Emotional Design: Why We Love (or Hate) Everyday Things*, stressing the role of aesthetics in our response to objects.

LOOKING BACK: CULTURES AND BRIDGES

Despite overlapping interests, in a dynamic environment with shifting alliances, the major threads of HCI research—HF&E, IS, LIS, and CHI—have not merged. They have interacted with each other only sporadically, although not for a lack of bridge-building efforts. The Human Factors Society co-organized the first CHI conference. CSCW sought to link CHI and IS. Mergers of OIS with CHI and later CSCW were considered. AIS SIGHCI tried to engage with CHI. Researchers recently hired into Information Schools remain active in the other fields.

Even within computer science, bridging is difficult. Researchers interested in interaction left SIGGRAPH to join the CHI community rather than form a bridge. A second opportunity arose 20 years later, when standard platforms powerful enough to support photorealism loomed, but the DUX conference series managed only three meetings. For AI, SIGART and SIGCHI cosponsor the Intelligent User Interface series, but participation has remained outside mainstream HCI. What are the obstacles to more extensive interaction across fields?

DISCRETION AS A MAJOR DIFFERENTIATOR

HF&E and IS arose before discretionary hands-on use was common. The information field only slowly distanced itself from supporting specialists. CHI occupied a new niche: discretionary use by nonexperts. HF&E and especially IS researchers considered organizational factors; CHI with few exceptions avoided domain-dependent work. As a consequence, HF&E and IS researchers shared journals. For example, Benbasat and Dexter (1985) published their work in *Management Science* and cited five *Human Factors* articles. Apart from LIS, they quickly focused on broad populations. IS countered its organizational focus by insisting that work be framed by theory, which set it apart from CHI in particular.

The appropriateness of a research method is tied to the motivation of the researchers. HF&E and CHI were shaped by psychologists trained in experimental testing of hypotheses about behavior, and hypothesis-driven experimentation was also embraced by IS. Experimental subjects agree to follow instructions for an extrinsic reward. This is a reasonable model for nondiscretionary use, but not for discretionary use. CHI researchers relabeled subjects as “participants,” which sounds volitional, and found that formal experimental studies were usually inappropriate: There were too many variables to test formally and feedback from a few participants was often enough. Laboratory studies of initial or casual discretionary use usually require confirmation in real-world settings anyway, more so than studies of expert or trained behavior, because of the artificial motivation of the laboratory study participant.

The same goals apply—fewer errors, faster performance, quicker learning, greater memorability, and being enjoyable—but the emphasis differs. For power plant operation, error reduction is critical, performance enhancement is good, and other goals are less important. For telephone order entry takers performance is critical, and testing an interface that could shave a few seconds from a repetitive operation requires a formal experiment. In contrast, consumers often respond to visceral appeal and initial experience. In assessing designs for mass markets, avoiding obvious problems can be more important than striving for an optimal solution. Less rigorous discount usability or cognitive walk-through methods (Nielsen 1989; Lewis et al. 1990) can be enough. Relatively time-consuming qualitative approaches, such as contextual design or persona use (Beyer and Holtzblatt 1998; Pruitt and Adlin 2006), can provide a deeper understanding when context is critical or new circumstances arise.

CHI largely abandoned its roots in scientific theory and engineering, which does not impress researchers from HF&E or theory-oriented IS. The controversial psychological method of verbal reports, developed by Newell and Simon (1972) and foreshadowed by gestalt psychology, was applied to design by Clayton Lewis as “thinking aloud” (Lewis and Mack 1982; Lewis 1983). Perhaps the most widely used CHI method, it led some researchers in other fields to characterize CHI people as wanting to talk about their experiences instead of doing research.

ACADEMIC, LINGUISTIC, AND GENERATIONAL CULTURES

The academic culture of the sciences is that conferences are venues for work in progress and journals are repositories for polished work. The disciplines of HF&E, IS, Documentation, and Library Science adhere to this practice. In contrast, for U.S. computer science disciplines, conference proceedings are now the final destination of most work. Outside the United States, computer science retains a journal focus, which suggests that a key factor was the decision of ACM to archive conference proceedings (Grudin 2010). Information science draws on researchers from both camps, journals as archival and conferences as archival. Of course, a difference in preferred channel impedes communication. Researchers in journal cultures chafe at CHI’s insistence on polish and its high conference rejection rates; CHI researchers are dismayed by the lack of polish at other conferences and are less inclined to read journals.

CHI conferences accept 20%–25% of submissions. With a few exceptions, HF&E and IS conferences accept about 50% or more. In contrast, CHI journals receive fewer submissions and have higher acceptance rates. Many CHI researchers report that journals are not relevant. By my estimate, at most 15% of the work in CHI-sponsored conferences reaches journal publication. In contrast, an IS track organizer for HICSS estimated that 80% of research there progressed to a journal (Jay Nunamaker, opening remarks at HICSS-38, January 2004).

A linguistic divide also set CHI apart. HF&E and IS use the term “operator” and a “user” could be a manager who read printed reports. For CHI, “operator” was demeaning and a “user” was always a hands-on user. In HF&E and IS streams, “task analysis” refers to an organizational decomposition of work, perhaps considering external factors; in CHI, “task analysis” is a cognitive decomposition, such as breaking a text editing move operation into select, cut, select, and paste. In IS “implementation” means organizational deployment, whereas in CHI it is a synonym for development. The terms “system,” “application,” and “evaluation” also have different connotations or denotations in the different fields. Significant misunderstandings resulted from failures to appreciate these differences.

Different perspectives and priorities were also reflected in attitudes toward standards. Many HF&E researchers contributed to standards development, believing that standards contribute to efficiency and innovation. A view widespread

in the CHI community was that standards inhibit innovation. Both views have elements of truth, and the positions partly converged as Internet and Web standards were tackled. However, the attitudes reflected the different demands of government contracting and commercial software development. Specifying adherence to standards is a useful tool for those preparing requests for proposals, whereas compliance with standards can make it more difficult for a product to differentiate itself.

The generational divide was also a factor. Many CHI researchers who grew up in the 1960s and 1970s did not appreciate the prior generation's orientation toward military, government, and business systems. They were also put off by the lack of gender neutrality in the HF&E and IS "man-machine interaction" literature, which one still occasionally encounters. Only in 1994 did *IJMMS* become *International Journal of Human-Computer Studies*. Such differences affected the enthusiasm for building bridges and exploring literatures.

Competition for resources was another factor. Computers of modest capability were extremely expensive for much of the time span we have considered. CHI was initially largely driven by the healthy tech industry, whereas research in the other fields was more dependent on government funding that waxed and waned. When funding waxed, demand for researchers outstripped supply. HCI prospered during AI winters, starting with Sutherland's use of the TX-2 when AI suffered its first setback and recurring with the emergence of major HCI laboratories during the severe AI winter of the late 1970s. Library schools laboring to create information science programs had to compete with computer science departments that awarded faculty positions to graduates of master's programs when the supply was low.

Greater interdisciplinarity is intellectually seductive. Could we not learn by looking over fences? But a better metaphor might be the big bang. Digital technology is an explosion, streaming matter and energy in every direction, forming worlds that at some later date might discover one another and find ways to communicate, and then again, might not.

LOOKING FORWARD: TRAJECTORIES

The future of HCI will be dynamic and full of surprises. The supralinear growth of hardware capability confounds efforts at prediction: We rarely experience exponential change and do not reason well about it. In the United States, NSF is tasked with envisioning the future and providing resources for taking us there, yet two major recent HCI initiatives, "Science of Design" and "CreativIT" (focused on creativity), wound down quickly. Nevertheless, extrapolations from observations about the past and present suggest possible developments, providing a prism through which to view other chapters in this handbook and perhaps some guidance in planning a career.

DISCRETION: NOW YOU SEE IT, NOW YOU DON'T

We exercise prerogative when we use digital technology—sometimes. More often when at home, less often at work. Sometimes we have no choice, as when confronted by a telephone answering system. Those who are young and healthy have more choices than those constrained by injury or aging.

Many technologies follow the maturation path shown in Figure 3. Software that was discretionary yesterday is indispensable today. Collaboration forces us to adopt shared conventions. Consider a hypothetical team that has worked together for 20 years. In 1990, members exchanged printed documents. One person still used a typewriter, whereas others used different word processors. One emphasized words by underlining, another by italicizing, and a third by bolding. In 2000, the group decided to exchange digital documents. They had to adopt the same word processor. Choice was curtailed; it was only exercised collectively. Today this team is happy sharing documents in PDF format, so they can again use different word processors. Perhaps tomorrow software will let them personalize their view of a single underlying document, so one person can again use and see in italics what another sees as bold or underlined.




FIGURE 3 From invention to maturity.

Shackel (1997, p. 981) noted this progression under the heading “From Systems Design to Interface Usability and Back Again.” Early designers focused at the system level; operators had to cope. When the PC merged the roles of operator, output user, and program provider, the focus shifted to the human interface and choice. Then individual users again became components in fully networked organizational systems. Discretion can evaporate when a technology becomes mission-critical, as word processing and e-mail did in the 1990s.

The converse also occurs. Discretion increases when employees can download free software, bring smartphones to work, and demand capabilities they enjoy at home. Managers are less likely to mandate the use of a technology that they use and find burdensome. For example, language-processing systems appealed to military officers, until they themselves became hands-on users:

Our military users ... generally flatly refuse to use any system that requires speech recognition. ... Over and over and over again, we were told “If we have to use speech, we will not take it. I don’t even want to waste my time talking to you if it requires speech. ...” I have seen generals come out of using, trying to use one of the speech-enabled systems looking really whipped. One really sad puppy, he said “OK, what’s your system like, do I have to use speech?” He looked at me plaintively. And when I said “No,” his face lit up, and he got so happy (Forbus 2003; see also Forbus, Usher, and Chapman [2003]).

In domains where specialized applications become essential and where security concerns curtail openness, discretion can recede. But Moore’s law (broadly construed), competition, and the ease of sharing bits should guarantee a steady flow of experimental technologies with unanticipated and thus initially discretionary uses.

UBIQUITOUS COMPUTING: INVISIBLE HUMAN–COMPUTER INTERACTION?

Norman (1988, p. 185) wrote of “the invisible computer of the future.” Like motors, he speculated, computers would be present everywhere and visible nowhere. We interact with clocks, refrigerators, and cars. Each has a motor, but who studies human–motor interaction? Marc Weiser subsequently introduced a similar concept, “ubiquitous computing.” A decade later, at the height of the Y2K crisis and the Internet bubble, computers were more visible than ever. But after a quarter century, while we may always want a large display or two, would anyone call a smartphone or a book reader a computer? The visions of Norman and Weiser may be materializing.

With digital technology embedded everywhere, concern with interaction is everywhere. HCI may become invisible through omnipresence. As interaction with digital technology becomes part of everyone’s research, the three long-standing HCI fields are losing participation.

Human Factors and Ergonomics

David Meister, author of *The History of Human Factors and Ergonomics* (1999), stresses the continuity of HF&E in the face of technology change:

Outside of a few significant events, like the organization of HFS in 1957 or the publication of Proceedings of the annual meetings in 1972, there are no seminal occurrences ... no sharp discontinuities that are memorable. A scientific discipline like HF has only an intellectual history; one would hope to find major paradigm changes in orientation toward our human performance phenomena, but there is none, largely because the emergence of HF did not involve major changes from pre-World War II applied psychology. In an intellectual history, one has to look for major changes in thinking, and I have not been able to discover any in HF (e-mail, September 7, 2004).

Membership in the Computer Systems Technical Group has declined. Technology is heavily stressed in technical groups such as Cognitive Engineering and Decision Making, Communication, Human Performance Modeling, Internet, System Development, and Virtual Environment. Nor do Aging, Medical Systems, or other technical groups avoid “invisible computers.”

Information Systems

While IS was thriving during the Y2K crisis and the Internet bubble, other management disciplines—finance, marketing, operations research, and organizational behavior—became more technically savvy. When the bubble burst and enrollments declined, the IS niche became less well defined. The research issues remain significant, but this cuts two ways. As IT organizations standardize on products and outsource IT functions, business-to-business and web portals for customers get more attention. These give rise to finance and marketing considerations, so HCI functions could be assumed by other management disciplines.

Computer–Human Interaction

This nomadic group started in psychology and then won a seat at the computer science table, which was bestowed grudgingly. Several senior CHI people moved to information schools. Lacking a well-defined academic niche, CHI ties its identity to the SIGCHI organization and the CHI conference. Membership in SIGCHI peaked in 1992 and conference attendance peaked in 2001. As new technologies become widely used, specialized conferences appear, often started by younger researchers. World Wide Web conferences included papers on HCI issues from the outset. HCI is an “invisible” presence in conferences on agents, design, and on computing that is ubiquitous, pervasive, accessible, social, and sustainable. High rejection rates for conference submissions and a new generational divide could accelerate the dispersion of research.

CHI attendance has become more exclusively academic, despite industry’s need for basic research in specific areas.

Apart from education and health, which have broad appeal, and software design and development, CHI remains largely focused on general phenomena and resistant to domain-specific work. This creates additional opportunities for regional and specialized conferences.

INFORMATION

Early in the computer era, there were no networks and memory was fantastically expensive. Computers were for computation, not information processing. Today, the situation is reversed: Memory and bandwidth are so plentiful that most computation is in the service of processing and distributing information. And the shift to an emphasis on information, with computation present but less visible, could well accelerate.

Cronin (1995) proposed that information access, in terms of intellectual, physical, social, economic, and spatial/temporal factors, is the focus of the information field. Information is acquired from sensors and human input; it flows over networks including the Web, and is aggregated, organized, and transformed. The routing and management of information within enterprises, as well as the consequences of ever more permeable organizational boundaries, is evolving. Approaches to personal information management are also rapidly changing. It was once centered on shoeboxes of photographs and boxes of old papers. Now most of us face significant online information management decisions, choosing what to keep locally, what to maintain in the cloud, and how to organize it to ensure its future accessibility. The CHI field has over a decade of work on information design and visualization (see Chapter 23 by Stuart Card).

In speculating about the future, Cronin (1995, p. 56) quotes Wersig (1992) who argued that concepts around information might function “like magnets or attractors, sucking the focus-oriented materials out of the disciplines and restructuring them within the information scientific framework.” Could this happen? Information schools have hired senior and junior people from many relevant areas. Andrew Dillon, dean of the University of Texas, School of Information, worked at Loughborough with Brian Shackel and Ken Eason. Syracuse, the first extant school of information (since 1974), has faculty with IS training and orientation. CHI faculty have migrated to information schools and departments of several leading universities.

Communication Studies is a discipline to watch. Rooted in humanities and social sciences, it is gradually assuming a quantitative focus. Centered on studies of television and other mass media, the field blossomed in the 1980s and 1990s. Only in the last several years has computer-mediated communication reached the scale of significance of other mass media. HCI is in a position to draw on past work in communication, as communication focuses more on digital media.

The rise of specialized programs—biomedical informatics, social informatics, community informatics, and information and communication technology for development

(ICT4D)—works against the consolidation of information studies. Information, like HCI, could become invisible through ubiquity. The annual Information Conference is a barometer. In 2005 and 2006, there was active discussion and disagreement about directions. Should new journals and research conferences be pursued, or should the field stick with the established venues in the various contributing disciplines? In the years since, faculty from different fields worked out pidgin languages with which to communicate with each other. Assistant professors were hired and graduate students enlisted, whose initial jobs and primary identities are with information. Will they creolize the pidgin language?

One can get a sense that the generals may still be arguing over directions, but the troops are starting to march. It is not clear where they will go. The generals, although busy with local campaigns, are reluctant to turn over command. The annual iConference vies with the less international but more established ASIST conference. However this evolves, in the long term, information is likely to be the major player in HCI. Design and information are active foci of HCI today, but the attention to design is compensation for past neglect. Information is being reinvented.

CONCLUSION: THE NEXT GENERATION

Looking back, cyclic patterns and cumulative influences are visible. New waves of hardware enabled different ways to support the same activities. E-mail arrived as an informal communication medium, was embraced by students, regarded with suspicion by organizations, and eventually became more formal and used everywhere. Then texting and instant messaging came along as an informal medium, were embraced by students, regarded with suspicion by organizations, and eventually became used everywhere. Social networking came along. ...

Mindful of Edgar Fiedler’s admonition that “he who lives by the crystal ball soon learns to eat ground glass,” consider this: In the mid-1980s, the mainframe market lost the spotlight. Organizations were buying hundreds of PCs, but these were weak devices with little memory, hard to network. They did not need more mainframes, but what about a massive, parallel supercomputer? Government and industry invested vast sums in high-performance computing only to discover that it was hard to decompose most computational problems into parallel processes whose output could be reassembled. As these expensive and largely ineffective efforts proceeded, PCs slowly got stronger, added some memory, got networked together, and, without vast expenditures and almost unnoticed at first, the Internet and the Web emerged.

Today the PC is losing the spotlight. Organizations buy hundreds of embedded systems, sensors, and effectors, but these are weak devices with little memory, hard to network. Some tasks can be handed off to a second processor, but how far can parallel multicore computers take us? Government and industry are investing large sums in parallel computing. They are rediscovering the difficulties. Sensors and effectors will add processing and memory, harvest energy,

and get networked. What will that lead to? The role of the PC may shift, becoming a personal control station where we can monitor vast quantities of information on anything of interest—our health, the state of household appliances, Internet activity, etc.—on large displays, with specific tasks easily moved to portable or distributed devices.

New technologies capture our attention, but of equal importance is the rapid maturation of technologies such as digital video and document repositories, as well as the complex specialization occurring in virtually all domains of application. Different patterns of use emerge in different cultures and different industries. Accessibility and sustainability are wide-open, specialized research and development areas. Tuning technologies for specific settings can bring human factors approaches to the fore; designing for efficient heavy use could revive command-driven interfaces, whether the commands are typed, spoken, or gestural.

Digital technology has inexorably increased the visibility of activity. We see people behaving not as we thought they would or as we think they should. Rules, conventions, policies, regulations, and laws are not consistently followed; sanctions for violating them are not uniformly applied. Privacy and our evolving attitudes toward it are a small piece of this powerful progression. Choosing how to approach these complex and intensifying challenges—Where do we increase enforcement? Should or could we create more nuanced rules? When should we tolerate more deviance?—at the levels of families, organizations, and societies. This will be a perpetual preoccupation as technology exposes the world as it is.

Until some time after it is revoked, Moore's law broadly construed will ensure that digital landscapes provide new forms of interaction to explore and new practices to improve. The first generation of computer researchers, designers, and users grew up without computers. The generation that followed used computers as students, entered workplaces, and changed the way technology was used. Now a generation has grown up with computers, game consoles, and cell phones. They absorbed an aesthetic of technology design and communicate by messaging. They are developing skills at searching, browsing, assessing, and synthesizing information. They use smartphones, acquire multimedia authoring talent, and embrace social networking sites. They have different takes on privacy and multitasking. They are entering workplaces, and everything will be changed once again. However it is defined and wherever it is studied, human-computer interaction will for some time be in its early days.

APPENDIX: PERSONAL OBSERVATIONS

My career from 1973 to 1993 followed a common enough path. I was one of many who worked as a computer programmer, studied cognitive psychology, spent time as an HCI professional in industry, and then moved to academia. I describe personal experiences here not because I am special, but to add texture and a sense of the human impact of some of the developments I have described. My interest in history arose from the feeling of being swept along by invisible forces,

sometimes against my intention. My first effort at understanding was titled "The Computer Reaches Out" (Grudin 1990): I saw computers evolving and slowly reaching into the world and changing it in ways that we, their developers, had not foreseen.

1970: A CHANGE IN PLANS

As a student, I read and believed the *Life* magazine article that forecast computers with superhuman intelligence arriving in several years. I concluded that if we survived a few years, we could count on machines to do all useful work. Human beings should focus on doing what they enjoy. I shifted from physics to mathematics and from politics to literature.

1973: THREE PROFESSIONS

Looking for my first job in 1973, I found three computer job categories in the *Boston Globe* classifieds: (1) operators, (2) programmers, and (3) systems analysts. Not qualified to be a highly paid analyst, I considered low-paid, hands-on operator jobs, but I landed a programming job with Wang Laboratories, which was at the time a small electronics company. For 2 years, I never saw the computer that my programs ran on. I flowcharted on paper and coded on coding sheets that a secretary sent to be punched and verified. A van carried the stack of cards 20 miles to a computer center, and later that day or the next day I got the printout. It might say something like "Error in Line 20," and I would resume work on the program.

1975: A CADRE OF DISCRETIONARY HAND-ON USERS

In 1975, Wang acquired a few teletype terminals with access to the WYLBUR line editor, developed at the Stanford Linear Accelerator. Some of us programmers chose to abandon paper and became hands-on computer users.

1983: CHILLY RECEPTION FOR A PAPER ON DISCRETION IN USE

My first HCI publication, Grudin and MacLean (1984), was written when I was a postdoctoral researcher at the MRC Applied Psychology Unit. Allan and I showed that people sometimes choose a slower interface for aesthetic or other reasons even when they are familiar with a more efficient alternative. A senior colleague asked us not to publish it. He worked on improving expert efficiency through cognitive modeling. A demonstration that greater efficiency could be undesirable would be a distraction, he said: "Sometimes the larger enterprise is more important than a small study."

1984: ENCOUNTERING MOORE'S LAW, INFORMATION SYSTEMS, HUMAN FACTORS, AND DESIGN

I returned to Wang, which had become a large minicomputer company, and found that Moore's law had changed

the business. More hardware was ordered from catalogs and the reduced cost of memory and other factors had changed programming priorities and skills. I was soon influenced by another cognitive psychologist, Susan Ehrlich, who worked in a marketing research group and later managed the human factors group. She introduced me to the IS literature, which I found difficult to understand. I attended Boston-area chapter meetings of both HFS and SIGCHI. I saw the cultural differences but felt CHI could learn from human factors. In a futile gesture to counter CHI antipathy toward human factors, I began calling myself a human factors engineer. I drove to Cambridge to see the newly released Macintosh. Few software engineers had the visual design skills that I realized would become important, so at work I looked for industrial designers of hardware (boxes) who could be enlisted to support software interface design.

1985: THE GRAPHICAL USER INTERFACE SHOCK

In the early 1980s, Phil Barnard and I were among the many cognitive psychologists working on command naming. This was an important application in the era of command-line interfaces, but the ambition was to develop a comprehensive theoretical foundation for HCI. The success of the Mac in 1985 curtailed interest in command names. No one would build on our past work—a depressing thought. It also dashed the hope for a comprehensive theoretical foundation for HCI. We had to choose: Am I a cognitive psychologist or a computer professional? Phil remained a psychologist.

1986: BEYOND THE USER: GROUPS AND ORGANIZATIONS

I agreed to join MCC, an industry research consortium. Between jobs I worked on two papers, each addressing a major challenge encountered in product development: (1) From 1984 to 1986, I had worked on several products or features intended to support groups rather than individual users. These had not done well. Why was group support so challenging? (2) It was painfully evident that organizational structures and development processes were badly suited to interactive software development. What could be done about it? These issues formed the basis for much of my subsequent research.

1989: DEVELOPMENT CONTEXTS: A MAJOR DIFFERENTIATOR

I spent 2 years at Aarhus University. Within weeks of arriving in a country that had little commercial software development, I saw that differences in the conditions that govern product, in-house, and contract development of interactive software could shape practices and perceptions in CHI, IS, and software engineering. Sorting this out led to my first library research for purely historical purposes (Grudin 1991). Perusing long-forgotten journals and magazines in dusty library corridors felt like wandering through an archaeological site.

1990: JUST WORDS: TERMINOLOGY CAN MATTER

I felt a premonition in 1987 when my IS-oriented colleague Susan Ehrlich titled a paper “Successful Implementation of Office Communication Systems.” By “implementation,” she meant introduction into organizations. To me, implementation was a synonym for coding or development. Sure enough, the ACM editor asked her to change the word implementation to adoption (Ehrlich 1987). What she called systems I called applications. Was language, usually an ally, getting in the way?

In 1990, I described the focus of my planned HCI course at Aarhus as “user-interface evaluation.” My new colleagues seemed embarrassed. Weeks later, a book written by one of them was published (Bødker 1990). Its first sentence was a quotation: “Design is where the action is, not evaluation.” Now I was embarrassed. In an in-house development world, with its dogma of getting the design right up front, development projects could take 10 years. Evaluation occurred at the end when only cosmetic changes were possible, and had a negative stigma. In commercial product development, evaluation of the previous version, competitive products, and (ideally) prototypes was integral to design. Evaluation is central to iterative design. It draws on the experimental psychologists’ skillset. We considered it a good thing.

Later in 1990, I participated in a panel on task analysis at a European conference. To my dismay, this IS-oriented group defined task analysis differently than I did. To them, it meant an organizational task analysis: tasks as components in a broad work process. In CHI, it meant a cognitive task analysis: breaking a simple task into components; for example, is “move text” thought of as “select-delete-paste” or as “select-move-place”? Some Europeans felt North American claims to have conducted task analyses were disgraceful, not understanding the context.

Also in 1990, en route to giving a job talk at the University of California Irvine, my first lecture to an IS audience at the University of California Los Angeles Anderson School of Management ended badly when the department head asked a question. It seemed meaningless, so I replied cautiously. He rephrased the question. I rephrased my response. He started again, then stopped and shrugged as if to say, “this fellow is hopeless.” When I saw him a few months later, he was astonished to learn that his Irvine friends were hiring me. Later, I understood the basis of our failure to communicate: We attached different meanings to the word “users.” To me, it meant hands-on computer users. He was asking about IS users who specified database requirements and read reports, but were not hands-on computer users. To me all use was hands-on, so his question had made no sense.

A book could be written about the word “user.” From a CHI perspective, the IS user was “the customer.” Consultants use “client.” In IS, the hands-on user was “the end-user.” In CHI parlance, end-user and user were one and the same—a person who entered data and used the output. The word end-user seemed superfluous or an affectation. Human factors

used “operator” which CHI considered demeaning. In software engineering, user typically denoted a tool user, that is, a software engineer.

A final terminology note: the male generic. I avoided submitting to *IJMMS* and turned down an invitation to speak at a “man–machine interaction” event. I was keen on learning from other disciplines, but that was a linguistic bridge I usually avoided crossing. I generally consider words to be a necessary but uninteresting medium for conveying meaning, but such experiences led to an essay on unintended consequences of language (Grudin 1993).

2010: REFLECTIONS ON BRIDGING EFFORTS

I have been a minor participant in efforts to find synergies drawing from CHI and HFS, OIS, IS (in both CSCW and AIS SIGHCI), or Design. None succeeded. I’ve interviewed others who participated years ago and identified the obstacles touched on in the introduction, many of which I experienced. As a boomer, I experienced generational and cultural divides. Many of my MCC colleagues joined the consortium to avoid Star Wars military projects. We lived through disputes between cognitive psychologists and radical behaviorists. I was among CHI researchers who shifted from journals to conferences as the primary publication venue and from hypothesis-driven experimentation to build-and-assess and qualitative field research.

Some differences fade over time, but many persist. Conference reviewers are often irritated by unfamiliar acronyms used by authors from other fields. Writing a chapter for an IS-oriented book (Palen and Grudin 2002), my coauthor and I wrangled at great length with the editor over terminology.

In researching this article, I reviewed the literature on TAM, the model of white-collar employee perceptions of technology that is heavily cited in IS but never in CHI. I unsuccessfully searched online for TAM references. Only on my third attempt did I see the problem: TAM stands for “Technology Acceptance Model,” but I repeatedly typed in “Technology Adoption Model.” TAM examined nondiscretionary acceptance, I think in terms of discretionary adoption. Different biases lead to different terminology, and confusion.

2010: Predicting the Future

Detailed forecasts, including mine, rarely look good upon close inspection. But understanding the forces that have shaped the past offers hope of anticipating or reacting quickly to future events. Even more useful may be indications of where effort will be futile. I believe the most common error is to underestimate the impact of hardware changes, and in particular that once effects start to be felt, how rapidly they will escalate. I published some analysis and projection in the November 2006 and January 2007 issues of *ACM Interactions*—check to see how I’m doing. (<http://interactions.acm.org/content/archives.php>).

REFERENCES*

- Ackoff, R. L. 1967. Management misinformation systems. *Manage Sci* 14:B147–56.
- Asimov, I. 1950. *I, Robot*. New York: Gnome Press.
- Aspray, W. 1999. Command and control, documentation, and library science: The origins of information science at the University of Pittsburgh. *IEEE Ann Hist Comput* 21(4):4–20.
- Baecker, R., and W. Buxton. 1987. A historical and intellectual perspective. In *Readings in HCI: A multidisciplinary approach*, ed. R. Baecker and W. Buxton, 41–54. San Francisco, CA: Morgan Kaufmann.
- Baecker, R., J. Grudin, W. Buxton, and S. Greenberg. 1995. A historical and intellectual perspective. In *Readings in HCI: Toward the Year 2000*, ed. R. Baecker, J. Grudin, W. Buxton and S. Greenberg, 35–47. San Francisco, CA: Morgan Kaufmann.
- Bagozzi, R. P., F. D. Davis, and P. R. Warshaw. 1992. Development and test of a theory of technological learning and usage. *Hum Relat* 45(7):660–86.
- Banker, R. D., and R. J. Kaufmann. 2004. The evolution of research on Information Systems: A fiftieth-year survey of the literature in Management Science. *Manage Sci* 50(3):281–98.
- Bannon, L. 1991. From human factors to human actors: The role of psychology and HCI studies in system design. In *Design at Work*, ed. J. Greenbaum and M. Kyng, 25–44. Hillsdale, NJ: Erlbaum.
- Bardini, T. 2000. *Bootstrapping: Douglas Engelbart, Coevolution, and the Origins of Personal Computing*. Stanford, CA: Stanford University.
- Barnard, P. 1991. Bridging between basic theories and the artifacts of HCI. In *Designing Interaction: Psychology at the Human-Computer Interface*, ed. J. M. Carroll, 103–27. Cambridge: Cambridge University Press.
- Barnard, P., J. May, D. Duke, and D. Duce. 2000. Systems, interactions, and macrotheory. *ACM Trans Comput Hum Interact* 7(2):222–62.
- Begeman, M., P. Cook, C. Ellis, M. Graf, G. Rein, and T. Smith. 1986. Project Nick: Meetings augmentation and analysis. In *Proceedings Computer-Supported Cooperative Work 1986*, 1–6. Austin, TX: CSCW ’86.
- Benbasat, I., and A. S. Dexter. 1985. An experimental evaluation of graphical and color-enhanced information presentation. *Manage Sci* 31(11):1348–64.
- Bennett, J. L. 1979. The commercial impact of usability in interactive systems. In *Man-computer communication*, ed. B. Shackel, Vol. 2, 1–17. Maidenhead, UK: Pergamon-Infotech.
- Bewley, W. L., T. L. Roberts, D. Schroit, and W. L. Verplank. 1983. Human factors testing in the design of Xerox’s 8010 “Star” office workstation. In *Proceedings CHI’83*, 72–77. New York: ACM.
- Beyer, H., and K. Holtzblatt. 1998. *Contextual Design—Defining Customer-Centered Systems*. San Francisco, CA: Morgan Kaufmann.
- Bjerknes, G., P. Ehn, and M. Kyng, eds. 1987. *Computers and Democracy—a Scandinavian Challenge*. Aldershot, UK: Avebury.
- Björn-Andersen, N., and B. Hedberg. 1977. Design of information systems in an organizational perspective. In *Prescriptive Models of Organizations*, ed. P. C. Nystrom and W. H. Starbuck, Vol. 5, 125–42. *TIMS Studies in the Management Sciences*. Amsterdam, Netherlands: North-Holland.

* All URLs were accessed December 9, 2011.

- Blackwell, A. 2006. The reification of metaphor as a design tool. *ACM Trans Comput Hum Interact* 13(4):490–530.
- Blythe, M. A., A. F. Monk, K. Overbeeke, and P. C. Wright, eds. 2003. *Funology: From Usability to User Enjoyment*. New York: Kluwer.
- Borman, L. 1996. SIGCHI: The early years. *SIGCHI Bull* 28(1):1–33. New York: ACM.
- Borman, L. ed. 1981. *Proceedings of the Joint Conference on Easier and More Productive Use of Computer Systems, Part II: Human interface and user interface*. New York: ACM.
- Bowman, W. J. 1968. *Graphic Communication*. New York: John Wiley.
- Buckland, M. 1998. Documentation, information science, and library science in the U.S.A. In *Historical Studies in Information Science*, ed. T. B. Hahn, and M. Buckland, 159–72. Medford, NJ: Information Today/ASIS.
- Buckland, M. 2009. As we may recall: Four forgotten pioneers. *Interactions* 16(6):76–69.
- Burke, C. 1994. *Information and Secrecy: Vannevar Bush, Ultra, and the Other Memex*. Lanham, MD: Scarecrow Press.
- Burke, C. 1998. A rough road to the information highway: Project INTREX. In *Historical studies in Information Science*, ed. T. B. Hahn and M. Buckland, 132–46. Medford, NJ: Information Today/ASIS.
- Burke, C. 2007. History of information science. In *Annual review of Information Science and Technology* 41, ed. B. Cronin, 3–53. Medford, NJ: Information Today/ASIST.
- Bush, V. 1945. As we may think. *Atl Mon* 176:101–8. <http://www.theatlantic.com/magazine/archive/1969/12/as-we-may-think/3881/>.
- Butler, P. 1933. *Introduction to Library Science*. Chicago, IL: Univ. of Chicago Press.
- Buxton, W. A. S. 2006. Early interactive graphics at MIT Lincoln Labs. <http://www.billbuxton.com/Lincoln.html>
- Bødker, S. 1990. *Through the Interface: A Human Activity Approach to User Interface Design*. Mahwah, NJ: Lawrence Erlbaum.
- Cakir, A., D. J. Hart, and T. F. M. Stewart. 1980. *Visual Display Terminals*. New York: Wiley.
- Card, S. K., and T. P. Moran. 1986. User technology: From pointing to pondering. In *Proceedings of the Conference on the History of Personal Workstations*, 183–98. New York: ACM.
- Card, S. K., T. P. Moran, and A. Newell. 1980a. Computer text-editing: An information-processing analysis of a routine cognitive skill. *Cogn Psychol* 12:396–410.
- Card, S. K., T. P. Moran, and A. Newell. 1980b. Keystroke-level model for user performance time with interactive systems. *Commun ACM* 23(7):396–410. New York: ACM.
- Card, S., T. P. Moran, and A. Newell. 1983. *The psychology of human-computer interaction*. Mahwah, NJ: Lawrence Erlbaum Associates.
- Carey, J. 1988. *Human Factors in Management Information Systems*. Greenwich, CT: Ablex.
- Carroll, J. M., ed. 2003. *HCI Models, Theories and Frameworks: Toward a Multidisciplinary Science*. San Francisco, CA: Morgan Kaufmann.
- Carroll, J. M., and R. L. Campbell. 1986. Softening up hard science: Response to Newell and Card. *Hum Comput Interact* 2(3):227–49.
- Carroll, J. M., and S. A. Mazur. 1986. Lisa learning. *IEEE Comput* 19(11):35–49.
- Cronin, B. 1995. Shibboleth and substance in North American Library and Information Science education. *Libri* 45:45–63.
- Damodaran, L., A. Simpson, and P. Wilson. 1980. *Designing Systems for People*. Manchester, UK: NCC Publications.
- Darrach, B. 1970. Meet Shaky: The first electronic person. *Life Mag* 69(21):58B–68.
- Davis, F. D. 1989. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Q* 13(3):319–39.
- Davis, G. B. 1974. *Management Information Systems: Conceptual Foundations, Structure, and Development*. New York: McGraw-Hill.
- Dennis, A., J. George, L. Jessup, J. Nunamaker, and D. Vogel. 1988. Information technology to support electronic meetings. *MIS Q* 12(4):591–624.
- DeSanctis, G., and R. B. Gallupe. 1987. A foundation for the study of group decision support systems. *Manage Sci* 33:589–610.
- Dyson, F. 1979. *Disturbing the Universe*. New York: Harper and Row.
- Ehrlich, S. F. 1987. Strategies for encouraging successful adoption of office communication systems. *ACM Trans Office Inf Syst* 5(4):340–57.
- Engelbart, D. 1962. Augmenting human intellect: A conceptual framework. SRI Summary Report AFOSR-3223. Reprinted in *Vistas in Information Handling*, ed. P. Howerton and D. Weeks, Vol. 1, 1–29. Washington, DC: Spartan Books, 1963. <http://www.doungengelbart.org/pubs/augment-3906.html>
- Engelien, B., and R. McBryde. 1991. *Natural Language Markets: Commercial Strategies*. London: Ovum Ltd.
- Evenson, S. 2005. Design and HCI highlights. Presented at the HCIC 2005 Conference, Winter Park, Colorado, February 6, 2005.
- Fano, R., and F. Corbato. 1966. Time-sharing on computers. *Sci Am* 214(9):129–40.
- Feigenbaum, E. A., and P. McCorduck. 1983. *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*. Reading, MA: Addison-Wesley.
- Fidel, R. 2011. *Human Information Interaction: An Ecological Approach to Information Behavior*. Cambridge, MA: MIT Press.
- Foley, J. D., and V. L. Wallace. 1974. The art of natural graphic man-machine conversation. *Proc IEEE* 62(4):462–71.
- Forbus, K. 2003. *Sketching for Knowledge Capture*. Lecture at Microsoft Research, Redmond, WA, May 2.
- Forbus, K. D., J. Usher, and V. Chapman. 2003. Qualitative spatial reasoning about sketch maps. In *Proceedings of the Innovative Applications of AI*, 85–92. Menlo Park: AAAI.
- Forster, E. M. 1909. *The Machine Stops*. Oxford and Cambridge Review, 8, November, 83–122.
- Friedman, A. 1989. *Computer Systems Development: History, Organization and Implementation*. New York: Wiley.
- Gilbreth, L. 1914. *The Psychology of Management: The Function of the Mind in Determining Teaching and Installing Methods of Least Waste*. NY: Sturgis and Walton.
- Good, I. J. 1965. Speculations concerning the first ultra-intelligent machine. *Adv Comput* 6:31–88. <http://commonsenseatheism.com/wp-content/uploads/2011/02/Good-Speculations-Concerning-the-First-Ultraintelligent-Machine.pdf>.
- Gould, J. D., and C. Lewis. 1983. Designing for usability—Key principles and what designers think. In *Proceedings of CHI '83*, 50–3. New York: ACM.
- Grandjean, E., and A. Vigliani. 1980. *Ergonomics Aspects of Visual Display Terminals*. London: Taylor and Francis.
- Gray, W. D., B. E. John, R. Stuart, D. Lawrence, and M. E. Atwood. 1990. GOMS meets the phone company: Analytic modeling applied to real-world problems. In *Proceedings of Interact '90*, 29–34. Amsterdam: North Holland.
- Greenbaum, J. 1979. *In the Name of Efficiency*. Philadelphia, PA: Temple University.
- Greif, I. 1985. Computer-Supported Cooperative Groups: What are the issues? In *Proceedings AFIPS Office Automation Conference*, 73–6. Montvale, NJ: AFIPS Press.

- Greif, I., ed. 1988. *Computer-Supported Cooperative Work: A Book of Readings*. San Mateo, CA: Morgan Kaufmann.
- Grudin, J. 1990. The computer reaches out: The historical continuity of interface design. In *Proceedings of CHI'90*, 261–8. New York: ACM.
- Grudin, J. 1991. Interactive systems: Bridging the gaps between developers and users. *IEEE Comput* 24(4):59–69.
- Grudin, J. 1993. Interface: An evolving concept. *Commun ACM* 36(4):110–9.
- Grudin, J. 2009. AI and HCI: Two fields divided by a common focus. *AI Mag* 30(4):48–57.
- Grudin, J. 2010. Conferences, community, and technology: Avoiding a crisis. In *Proceedings iConference 2010*. <https://www.ideals.illinois.edu/handle/2142/14921>.
- Grudin, J. 2011. Human-computer interaction. In *Annual Review of Information Science and Technology*, ed. B. Cronin, Vol. 45, 369–430. Medford, NJ: Information Today (for ASIST).
- Grudin, J., and A. MacLean. 1984. Adapting a psychophysical method to measure performance and preference tradeoffs in human-computer interaction. In *Proceedings of INTERACT'84*, 338–42. Amsterdam, Netherlands: North Holland.
- Hertzfeld, A. 2005. *Revolution in the Valley: The Insanely Great Story of How the Mac Was Made*. Sebastopol, CA: O'Reilly Media.
- HFES. 2010. HFES history. In *HFES 2010–2011, Directory and Yearbook*, 1–3. Santa Monica, CA: Human Factors and Ergonomics Society. Also found at <http://www.hfes.org/web/AboutHFES/history.html>
- Hiltz, S. R., and M. Turoff. 1978. *The Network Nation*. Reading, MA: Addison-Wesley.
- Hiltzik, M. A. 1999. *Dealers of Lightning: Xerox PARC and the Dawn of the Computer Age*. New York: HarperCollins.
- Hopper, G. 1952. The education of a computer. In *Proceedings of ACM Conference*, Reprinted in *Annals of the History of Computing*, 9(3–4), 271–81, 1987.
- Huber, G. 1983. Cognitive style as a basis for MIS and DSS designs: Much ado about nothing? *Manage Sci* 29(5):567–79.
- Hutchins, E. L., J. D. Hollan, and D. A. Norman. 1986. Direct manipulation interfaces. In *User Centered System Design*, ed. D. A. Norman and S. W. Draper, 87–124. Mahwah, NJ: Lawrence Erlbaum.
- Israelski, E., and A. M. Lund. 2003. The evolution of HCI during the telecommunications revolution. In *The Human-Computer Interaction Handbook*, ed. J. A. Jacko and A. Sears, 772–89. Mahwah, NJ: Lawrence Erlbaum.
- Johnson, T. 1985. *Natural Language Computing: The Commercial Applications*. London, UK: Ovum Ltd.
- Kao, E. 1998. *The History of AI*. <http://www.generation5.org/content/1999/aihistory.asp> (accessed March 13, 2007).
- Kay, A., and A. Goldberg. 1977. Personal dynamic media. *IEEE Comput* 10(3):31–42.
- Keen, P. G. W. 1980. MIS research: Reference disciplines and a cumulative tradition. In *First International Conference on Information Systems*, 9–18. Chicago, IL: Society for Management Information Systems.
- Kling, R. 1980. Social analyses of computing: Theoretical perspectives in recent empirical research. *Comput Surv* 12(1):61–110.
- Landau, R., J. Bair, and J. Siegmna, eds. 1982. Emerging office systems. In *Extended Proceedings of the 1980 Stanford International Symposium on Office Automation*. Norwood, NJ.
- Lenat, D. 1989. When will machines learn? *Mach Learn* 4:255–7.
- Lewis, C. 1983. The ‘thinking aloud’ method in interface evaluation. Tutorial given at CHI'83. Unpublished notes.
- Lewis, C., and R. Mack. 1982. Learning to use a text processing system: Evidence from “thinking aloud” protocols. In *Proceedings of the Conference on Human Factors in Computing Systems*, 387–92. New York: ACM.
- Lewis, C., P. Polson, C. Wharton, and J. Rieman. 1990. Testing a walk-through methodology for theory-based design of walk-up-and-use Interfaces. In *Proceedings of the Conference on Human Factors in Computing Systems*, 235–42. New York: ACM.
- Licklider, J. C. R. 1960. Man-computer symbiosis. *IRE Trans Hum Factors Electron HFE-1* 1:4–11. <http://groups.csail.mit.edu/medg/people/psz/Licklider.html>.
- Licklider, J. C. R. 1963. MEMORANDUM FOR: Members and Affiliates of the Intergalactic Computer Network. April 23. <http://www.kurzweilai.net/memorandum-for-members-and-affiliates-of-the-intergalactic-computer-network>.
- Licklider, J. C. R. 1965. *Libraries of the Future*. Cambridge, MA: MIT Press.
- Licklider, J. C. R. 1976. User-oriented interactive computer graphics. In *Proceedings of SIGGRAPH Workshop on User-Oriented Design of Interactive Graphics Systems*, 89–96. New York: ACM.
- Licklider, J. C. R., and W. Clark. 1962. On-line man-computer communication. *AFIPS Conf Proc* 21:113–28.
- Lighthill, J. 1973. Artificial intelligence: A general survey. In *Artificial Intelligence: A Paper Symposium*, ed. J. Lighthill, N. S. Sutherland, R. M. Needham, H. C. Longuet-Higgins, and D. Michie. London, UK: Science Research Council of Great Britain. http://www.chilton-computing.org.uk/inf/literature/reports/lighthill_report/p001.htm
- Long, J. 1989. Cognitive ergonomics and human-computer interaction. In *Cognitive Ergonomics and Human-Computer Interaction*, ed. J. Long and A. Whitefield, 4–34. Cambridge: Cambridge University Press.
- Machlup, F., and U. Mansfield, eds. 1983. *The Study of Information: Interdisciplinary Messages*. New York: Wiley.
- March, A. 1994. Usability: the new dimension of product design. *Harv Bus Rev* 72(5):144–9.
- Marcus, A. 2004. Branding 101. *ACM Interact* 11(5):14–21.
- Markoff, J. 2005. *What the Dormouse Said: How the 60s Counter-Culture Shaped the Personal Computer*. London, UK: Viking.
- Markus, M. L. 1983. Power, politics, and MIS implementation. *Commun ACM* 26(6):430–44.
- Martin, J. 1973. *Design of Man-Computer Dialogues*. New York: Prentice-Hall.
- McCarthy, J. 1960. Functions of symbolic expressions and their computation by machine, part 1. *Commun ACM* 3(4):184–95.
- McCarthy, J. 1988. B. P. Bloomfield, The question of artificial intelligence: Philosophical and sociological perspectives. *Ann Hist Comput* 10(3):224–9. <http://www-formal.stanford.edu/jmc/reviews/bloomfield/bloomfield.html>.
- Meister, D. 1999. *The History of Human Factors and Ergonomics*. Mahwah, NJ: Lawrence Erlbaum.
- Moggridge, B. 2007. *Designing Interactions*. Cambridge: MIT Press.
- Moravec, H. 1998. When will computer hardware match the human brain? *J Evol Technol* 1:1. <http://www.transhumanist.com/volume1/moravec.htm>.
- Mumford, E. 1971. A comprehensive method for handling the human problems of computer introduction. *IFIP Congr* 2:918–23.
- Mumford, E. 1976. Toward the democratic design of work systems. *Personnel Manage* 8(9):32–5.
- Myers, B. A. 1998. A brief history of human computer interaction technology. *ACM Interact* 5(2):44–54.
- National Science Foundation. 1993. *NSF 93–2: Interactive Systems Program Description*. <http://www.nsf.gov/pubs/stis1993/nsf932/nsf932.txt>

- National Science Foundation. 2003. NSF Committee of Visitors Report: Information and Intelligent Systems Division, 28 July 2003.
- Negroponte, N. 1970. *The Architecture Machine: Towards a More Humane Environment*. Cambridge: MIT Press.
- Nelson, T. 1965. A file structure for the complex, the changing, and the indeterminate. In *Proceedings of the ACM National Conference*, 84–100. New York: ACM.
- Nelson, T. 1973. A conceptual framework for man-machine everything. In *Proceedings of the National Computer Conference*, M21–6. Montvale, NJ: AFIPS Press.
- Nelson, T. 1996. Generalized links, micropayment and transcopy-right. <http://www.almaden.ibm.com/almaden/npuc97/1996/tnelson.htm>.
- Newell, A., and S. K. Card. 1985. The prospects for psychological science in human-computer interaction. *Hum Comput Interact* 1(3):209–42.
- Newell, A., and H. A. Simon. 1956. The logic theory machine: A complex information processing system. *IRE Trans Inf Theory* IT-2:61–79.
- Newell, A., and H. A. Simon. 1972. *Human Problem Solving*. New York: Prentice-Hall.
- Newman, W. M., and R. F. Sproull. 1973. *Principles of Interactive Computer Graphics*. New York: McGraw-Hill.
- Nielsen, J. 1989. Usability engineering at a discount. In *Designing and Using Human-Computer Interfaces and Knowledge Based Systems*, ed. G. Salvendy and M. J. Smith, 394–401. Amsterdam: Elsevier.
- Nielsen, J., and R. Molich. 1990. Heuristic evaluation of user interfaces. In *Proceedings of CHI'90*, 249–56. New York: ACM.
- Norberg, A. L., and J. E. O'Neill. 1996. *Transforming Computer Technology: Information Processing for the Pentagon*, 1962–86. Baltimore, MD: Johns Hopkins.
- Norman, D. A. 1982. Steps toward a cognitive engineering: Design rules based on analyses of human error. In *Proceedings of the Conference on Human Factors in Computing Systems*, 378–82. New York: ACM.
- Norman, D. A. 1983. Design principles for human-computer interfaces. In *Proc. CHI'83*, 1–10. New York: ACM.
- Norman, D. A. 1986. Cognitive engineering. In *User Centered System Design*, ed. D. A. Norman and S. W. Draper, 31–61. Mahwah, NJ: Lawrence Erlbaum.
- Norman, D. A. 1988. *Psychology of Everyday Things. Reissued as 'Design of Everyday Things' in 1990*. New York: Basic Books.
- Norman, D. A. 2004. *Emotional Design: Why We Love (or hate) Everyday Things*. New York: Basic Books.
- Nunamaker, J., R. O. Briggs, D. D. Mittleman, D. R. Vogel, and P. A. Balthazard. 1997. Lessons from a dozen years of group support systems research: A discussion of lab and field findings. *J Manage Inf Syst* 13(3):163–207.
- Nygaard, K. 1977. Trade union participation. Presentation at CREST Conference on Management Information Systems, Stafford, UK.
- Oakley, B. W. 1990. Intelligent knowledge-based systems—AI in the U.K. In *The Age of Intelligent Machines*, ed. R. Kurzweil, 346–9. Cambridge, MA: MIT Press.
- Palen, L., and J. Grudin. 2002. Discretionary adoption of group support software. In *Implementing Collaboration Technology in Industry*, ed. B. E. Munkvold, 159–90. London, UK: Springer-Verlag.
- Perlman, G., G. K. Green, and M. S. Wogalter. 1995. *Human Factors Perspectives on Human-Computer Interaction*. Santa Monica: Human Factors and Ergonomics Society.
- Pew, R. 2003. Evolution of HCI: From MEMEX to Bluetooth and beyond. In *The Human-Computer Interaction Handbook*, ed. J. A. Jacko and A. Sears, 1–17. Mahwah, NJ: Lawrence Erlbaum.
- Pruitt, J., and T. Adlin. 2006. *The Persona Lifecycle: Keeping People in Mind Throughout Product Design*. San Francisco, CA: Morgan Kaufmann.
- Rasmussen, J. 1980. The human as a system component. In *Human Interaction with Computers*, ed. H. T. Smith and T. R. G. Green, 67–96. London, UK: Academic.
- Rasmussen, J. 1986. *Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering*. New York: North-Holland.
- Rayward, W. B. 1983. Library and information sciences: Disciplinary differentiation, competition, and convergence. In *The Study of Information: Interdisciplinary Messages*, ed. F. Machlup and U. Mansfield, 343–405. New York: Wiley.
- Rayward, W. B. 1998. The history and historiography of Information Science: Some reflections. In *Historical Studies in Information Science*, ed. T. B. Hahn, and M. Buckland, 7–21. Medford, NJ: Information Today/ASIS.
- Remus, W. 1984. An empirical evaluation of the impact of graphical and tabular presentations on decision-making. *Manage Sci* 30(5):533–42.
- Roscoe, S. N. 1997. *The Adolescence of Engineering Psychology*. Santa Monica, CA: Human Factors and Ergonomics Society.
- Sammet, J. 1992. Farewell to Grace Hopper—End of an era! *Commun ACM* 35(4):128–31.
- Shackel, B. 1959. Ergonomics for a computer. *Design* 120:36–9.
- Shackel, B. 1962. Ergonomics in the design of a large digital computer console. *Ergonomics* 5:229–41.
- Shackel, B. 1997. HCI: Whence and whither? *J ASIS* 48(11):970–86.
- Shannon, C. E. 1950. Programming a computer for playing chess. *Philos mag* 7(41):256–75.
- Shannon, C. E., and W. Weaver. 1949. *The Mathematical Theory of Communication*. Urbana, IL: Univ. of Illinois Press.
- Sheil, B. A. 1981. The psychological study of programming. *ACM Comput Surv* 13(1):101–20.
- Shneiderman, B. 1980. *Software Psychology: Human Factors in Computer and Information Systems*. Cambridge, MA: Winthrop.
- Simon, H. A. 1960. *The New Science of Management Decision*. New York: Harper.
- Smith, H. T., and T. R. G. Green, eds. 1980. *Human Interaction with Computers*. Orlando, FL: Academic.
- Smith, S. L. 1963. Man-computer information transfer. In *Electronic Information Display Systems*, ed. J. H. Howard, 284–99. Washington, DC: Spartan Books.
- Smith, S. L., B. B. Farquhar, and D. W. Thomas. 1965. Color coding in formatted displays. *J Appl Psychol* 49:393–8.
- Smith, S. L., and N. C. Goodwin. 1970. Computer-generated speech and man-computer interaction. *Hum Factors* 12:215–23.
- Smith, S. L., and J. N. Mosier. 1986. *Guidelines for Designing User Interface Software (ESD-TR-86-278)*. Bedford, MA: MITRE.
- Suchman, L. 1987a. *Plans and Situated Action: The Problem of Human-Machine Communication*. Cambridge, UK: Cambridge University Press.
- Suchman, L. 1987b. Designing with the user: Review of 'Computers and democracy: A Scandinavian challenge. *ACM TOIS* 6(2):173–83.
- Sutherland, I. 1963. *Sketchpad: A Man-Machine Graphical Communication System*. Doctoral Dissertation, MIT. <http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-574.pdf>

- Taylor, F. W. 1911. *The Principles of Scientific Management*. New York: Harper.
- Turing, A. 1949. *Letter in London Times*, June 11. See Highlights from the Computer Museum report Vol. 20, Summer/Fall 1987. <http://ed-thelen.org/comp-hist/TCMR-V20.pdf>.
- Turing, A. 1950. Computing machinery and intelligence. *Mind* 49:433–60. Republished as “Can a machine think?” in J. R. Newman (Ed.), *The World of Mathematics*, (Vol. 4, pp. 2099–2123). New York: Simon and Schuster.
- Vessey, I., and D. Galletta. 1991. Cognitive fit: An empirical test of information acquisition. *Inf Syst Res* 2(1):63–84.
- Waldrop, M. M. 2001. *The Dream Machine: J.C.R. Licklider and the Revolution that Made Computing Personal*. New York: Viking.
- Weinberg, G. 1971. *The Psychology of Computer Programming*. New York: Van Nostrand Reinhold.
- Wells, H. G. 1905. *A Modern Utopia*. London, UK: Jonathan Cape. <http://www.gutenberg.org/etext/6424>.
- Wells, H. G. 1938. *World Brain*. London, UK: Methuen.
- Wersig, G. 1992. Information science and theory: A weaver bird’s perspective. In *Conceptions of Library and Information Science: Historical, Empirical, and Theoretical Perspectives*, ed. P. Vakkari and B. Cronin, 201–17. London, UK: Taylor Graham.
- White, P. T. 1970. Behold the computer revolution. *National Geographic* 38(5): 593–633. <http://blog.modernmechanix.com/2008/12/22/behold-the-computer-revolution/>
- Yates, J. 1989. *Control Through Communication: The Rise of System in American Management*. Baltimore, MD: Johns Hopkins.
- Zhang, P. 2004. AIS SIGHCI three-year report. *SIGHCI newslett* 3(1):2–6.
- Zhang, P., F. F. -H. Nah, and J. Preece. 2004. HCI studies in management information systems. *Behav Inf Technol* 23(3):147–51.
- Zhang, P., Li, N., Scialdone, M.J. & Carey, J. 2009. The intellectual advancement of Human-Computer Interaction Research: A critical assessment of the MIS Literature. *Trans Human-Computer Interaction*, 1(3): 55–107.

Part I

Humans in HCI

1 Perceptual-Motor Interaction

Some Implications for Human–Computer Interaction

*Timothy N. Welsh, Sanjay Chandrasekharan, Matthew Ray,
Heather Neyedli, Romeo Chua, and Daniel J. Weeks*

CONTENTS

1.1	Perceptual-Motor Interaction: A Behavioral Emphasis.....	3
1.1.1	Human Information Processing and Perceptual-Motor Behavior.....	4
1.1.2	Sensory Information during the Planning and Control of Action.....	5
1.1.3	Translation, Coding, and Mapping	6
1.2	Perceptual-Motor Interaction: Attention and Performance.....	7
1.2.1	Attention	7
1.2.1.1	Characteristics of Attention.....	7
1.2.1.2	Shifts of Attention.....	8
1.2.1.3	Action-Centered Attention.....	9
1.2.1.4	Summary.....	12
1.3	Common Coding Accounts of Perceptual-Motor Interactions.....	12
1.3.1	Perception–Action Common Coding.....	13
1.3.2	Imagination–Action Common Coding.....	14
1.3.3	Summary	14
1.4	Perceptual-Motor Interaction in Applied Tasks: A Few Examples.....	15
1.4.1	Attention Cueing for Military Target Detection.....	15
1.4.2	Deriving Novel Interaction Designs from Common Coding.....	15
1.5	Summary	16
	Acknowledgments.....	17
	References.....	17

1.1 PERCEPTUAL-MOTOR INTERACTION: A BEHAVIORAL EMPHASIS

Human–computer interaction is going through a period of rapid evolution. Although mouse, keyboard, and joystick devices will continue to dominate for the immediate future, embodied, gestural, and tangible interfaces—where individuals use their body to directly manipulate information objects—are rapidly changing the computing landscape. Most new laptops and mobile devices now support multitouch, which allows us to use our fingers and gestures to directly manipulate virtual objects on the screen. Hence, as an alternative to pointing and clicking with a mouse, we can now directly pull, push, grab, pinch, squeeze, crush, and throw virtual objects. We can shake our portable music player to change the song we are listening to, or we can turn our mobile devices horizontally to get a wider display screen. Using the Wii Remote, we can now use our body movements to interact with objects in video games and manipulate them. Tangible and augmented interfaces now allow us to interact

directly with virtual environments by moving actual objects on a tabletop (Hornecker et al. 2008). In sum, instead of being forced to use dissociated (mouse) and/or arbitrary (keyboard and joystick) sensorimotor mappings to achieve our goals, these new modes of interaction allow for a more direct mapping of our movements on to the work space. The “naturalness” and ease of operation of these interfaces are, in large part, due to the sensory and motor systems’ close connection to cognition. Therefore, as these new interfaces become more popular, it is becoming increasingly important to consider the mechanisms that support such interactions.

In our studies of human–computer interaction (HCI) and perceptual-motor interactions in general, we have adopted a number of theoretical and analytical frameworks as part of an integrated approach. Our chapters in earlier editions of this handbook (Chua, Weeks, and Goodman 2003; Welsh et al. 2007) reviewed much of this research and its implications for HCI. The emphasis for these earlier chapters was on using information-processing approaches to understand the translation of perceptual into motor space and the interaction

between processes of attention and action planning. Although our research has continued to explore the interplay between the processes of action and attention that we introduced in our chapter in the second edition (Welsh et al. 2007), we are intrigued by the possibilities offered by recently developed tangible interfaces and how theories of embodied cognition and common coding can support and enhance the progress of these systems.

Thus, in the present chapter, we have provided an updated review and expansion of our recent work in the area of action-centered attention and suggest some important implications for the role that action planning plays in the capture of attention and perception. We believe this work has important implications for the design of interaction modes. In the second section, we review the critical features of an alternative theoretical approach to cognition and action that presupposes an in-depth interaction between perception, cognition, and action. This latter theory has shaped much of our more recent work on the development of a tangible and embodied HCI. The critical theme that binds the seemingly diverse lines of work is the role that action planning has in information-processing systems. It is this central consideration that we argue has been lacking over the years and should be an important consideration for the future work.

1.1.1 HUMAN INFORMATION PROCESSING AND PERCEPTUAL-MOTOR BEHAVIOR

The information-processing framework has traditionally provided a major theoretical and empirical platform for many scientists interested in perceptual-motor behavior. The study of perceptual-motor behavior within this framework has inquired into such issues as the information capacity of the motor system (e.g., Fitts 1954), the attentional demands of movements (e.g., Posner and Keele 1969), motor memory (e.g., Adams and Dijkstra 1966), and processes of motor learning (e.g., Adams 1971). The language of information processing (e.g., Broadbent 1958) has provided the vehicle for discussions of mental and computational operations of the cognitive and perceptual-motor system (Posner 1982). Of interest in the study of perceptual-motor behavior is the nature of the cognitive processes that underlie perception and action.

The information-processing approach describes the human as an active processor of information, in terms that are now commonly used to describe complex computing mechanisms. An information-processing analysis describes observed behavior in terms of the encoding of perceptual information, the manner in which internal psychological subsystems utilize the encoded information, and the functional organization of these subsystems. At the heart of the human cognitive system are processes of information transmission, translation, reduction, collation, storage, and retrieval (e.g., Fitts 1964; Marteniuk 1976; Stelmach 1982; Welford 1968). Consistent with a general model of human information processing (e.g., Fitts and Posner 1967), three basic processes

have been distinguished historically. For our purposes, we refer to these processes as stimulus identification, response selection, and response programming. Briefly, stimulus identification is associated with processes responsible for the perception of information. Response selection pertains to the translation between stimuli and responses and the selection of a response. Response programming is associated with the organization of the final output (see Proctor and Vu 2003 or the present volume).

A key feature of early models of information processing is the emphasis upon the cognitive activities that precede action (Marteniuk 1976; Stelmach 1982). From this perspective, action is viewed only as the end-result of a complex chain of information-processing activities (Marteniuk 1976). Thus, chronometric measures such as reaction time and movement time, as well as other global outcome measures, are often the predominant dependent measures. However, even a cursory examination of the literature indicates that the time to engage a target has been a primary measure of interest. For example, a classic assessment of perceptual-motor behavior in the context of HCI and input devices was conducted by Card et al. (1978); see also English, Engelhart, and Berman (1967). Employing measures of error and speed, Card et al. (1978) had subjects complete a cursor-positioning task using four different control devices (mouse, joystick, step keys, and text keys). The data revealed the now well-known advantage for the mouse. Of interest is that the speed measure was decomposed into “homing” time, the time that it took to engage the control device and initiate cursor movement, and “positioning” time, the time to complete the cursor movement. Although the mouse was actually the poorest device in terms of the homing time measure, the advantage in positioning time produced the faster overall time. That these researchers sought to glean more information from the time measure acknowledges the importance of the movement itself in perceptual-motor interactions such as these.

The fact that various pointing devices depend on hand movement to control cursory movement has led researchers in HCI to emphasize Fitts’s law (Fitts 1954) as a predictive model of time to engage a target. The law predicts pointing (movement) time as a function of the distance to and the width of the target—where, in order to maintain a given level of accuracy, movement time must increase as the distance of the movement increases and/or the width of the target decreases. The impact of Fitts’s law is most evident by its inclusion in the battery of tests to evaluate computer pointing devices in ISO 9241-9. We argue that there are a number of important limitations to an exclusive reliance on Fitts’s law in this context.

First, although the law predicts movement time, it does so on the basis of distance and target size. Consequently, it does not allow for determining what other factors may influence movement time. Specifically, Fitts’s law is often based on a movement to a single target at any given time (although it was originally developed using reciprocal movements between two targets). However, in most HCI and graphical

user interface contexts, there is an array of potential targets that can be engaged by an operator. These nontarget, but action-relevant stimuli in the movement environment can have profound and unexpected effects on action planning and execution. For example, Adam et al. (2006) have repeatedly found that the last target in an array enjoys a movement time advantage that is not predicted by Fitts's law. In contrast, distracting nontarget stimuli that capture attention can negatively affect both the temporal and physical characteristics of the movements to the imperative target. We will discuss these negative consequences in greater detail in Section 1.2.1.3.

Second, we suggest that the emphasis on Fitts's law has diverted attention from the fact that cognitive processes involving the selection of a potential target from an array are an important, and time-consuming, information-processing activity that must precede movement to that target. For example, the Hick-Hyman law (Hick 1952; Hyman 1953) predicts the decision time required to select a target response from a set of potential responses—where the amount of time required to choose the correct response increases with the number of possible alternative responses. What is important to understand is that the two laws work independently to determine the total time it takes for an operator to acquire the desired location. In one instance, an operator may choose to complete the decision making and movement components sequentially. Under these conditions, the total time to complete the task will be the sum of the times predicted by the Hick-Hyman and Fitts's laws. Alternatively, an operator may opt to make a general movement that is an approximate average of the possible responses and then select the final target destination while the movement is being completed. Under such conditions, Hoffman and Lim (1997) reported interference between the decision and movement component that was dependent on their respective difficulties (see also Meegan and Tipper 1998).

Finally, although Fitts's law predicts movement time given a set of movement parameters, it does not actually reveal much about the underlying movement itself. Indeed, considerable research effort has been directed toward revealing the movement processes that give rise to Fitts's law. For example, theoretical models of limb control have been forwarded that propose that Fitts's law emerges as a result of multiple sub-movements (e.g., Crossman and Goodeve 1963/1983), or as a function of both initial movement impulse variability and subsequent corrective processes late in the movement (Meyer et al. 1988). These models highlight the importance of conducting detailed examinations of movements themselves as a necessary complement to chronometric explorations.

For these reasons, HCI situations that involve dynamic perceptual-motor interactions may not be best indexed merely by chronometric methods (cf. Card et al. 1978). Indeed, as HCI moves beyond the simple key press interfaces that are characteristic of early systems to include virtual and augmented reality, teleoperation, gestural and haptic interfaces, among others, the dynamic nature of perceptual-motor interactions are even more evident. Consequently, assessment

of the actual movement required to engage such interfaces will be more revealing.

To supplement chronometric explorations of basic perceptual-motor interactions, motor behavior researchers have also advocated a "movement process" approach (Kelso 1982). The argument is that in order to understand the nature of movement organization and control, analyses should also encompass the movement itself, and not just the activities preceding it (e.g., Kelso 1982, 1995; Marteniuk, MacKenzie, and Leavitt 1988). Thus, investigators have examined the kinematics of movements in attempts to further understand the underlying organization involved (e.g., Brooks 1974; Chua and Elliott 1993; Elliott et al. 1991; Kelso, Southard, and Goodman 1979; MacKenzie et al. 1987; Marteniuk et al. 1987). The relevance of this approach will become apparent in later sections.

1.1.2 SENSORY INFORMATION DURING THE PLANNING AND CONTROL OF ACTION

It almost goes without saying that different types of actions need different types and amounts of information to ensure accurate completion. Theoretical and experimental considerations of this issue, in a manner that is relevant to the field of HCI, have been expanded recently. Before discussing the evidence supporting this view and outlining some potential implications for HCI, we will briefly review the processes involved in the planning and control of action and the types of information used during these processes. Readers interested in gaining a more in-depth understanding of this research should consult a recent book on the topic (Elliott and Khan 2010).

Since the seminal work of Woodworth (1899), it has been generally accepted that goal-directed action consists of two main components: (1) the ballistic or open-loop component that initiates the action toward the goal and (2) the current control or closed-loop component during which movement-produced information is used to facilitate movement accuracy. The initial open-loop component is thought to represent the results of the stages of information processing and initial plan or motor program the individual has developed to complete the goal successfully. The second component of the action begins after the movement has been initiated and directed toward the goal. During this part of the movement, sources of movement-produced information about the current location and trajectory of the effector (feedback) are compared with the predicted or desired location and trajectory to determine any differences between the actual and desired movement pattern (i.e., movement error). These error signals are then used to correct the unfolding movement and achieve the goal.

The main evidence in favor of the notion of planning and control components for goal-directed actions is derived from detailed analyses of the kinematic profiles of aiming actions performed under various stimulus conditions (e.g., Chua and Elliott 1993; Heath 2005; see Khan et al. 2006 for a review).

Although vestibular and proprioceptive information is also necessary for the accurate planning and control of movement, visual information is by far the dominant source and, as such, is the source of information that is most commonly manipulated in these studies. As one would expect, people are more accurate and less variable under conditions in which they have vision of the environment than when they do not. The increase in accuracy is thought to occur, in large part, because the actor has visual information of both the effector and the target to detect and correct errors in the trajectory.

Of greater importance to the present discussion, however, are the results of the in-depth kinematic analysis of the aiming movements. The consistent finding of this research is that the majority of the differences between the movements executed with and without vision appears in the later portions of the movement. Specifically, the initial segments of movements performed both with and without vision are characterized by relatively similar smooth increases and then decreases in velocity (bell-shaped profiles). It is thought that this relative consistency arises because there is a relative consistency in the motor programs that are the basis of these early portions of the movement in both vision and no vision conditions. In contrast to the similarities in the initial parts of the movements, the later portions of the movement performed with continuous visual information of the environment are characterized by a much larger number of sudden decelerations and reaccelerations than movements executed in the absence of visual information. These discontinuities in the kinematic profiles are thought to represent instances in which the actor has used visual information about the effector and the target to detect errors and then formulate and execute corrective submovements. These online corrections increase the accuracy of the movement. In the absence of vision, most errors go undetected leading to smoother deceleration phases (i.e., with fewer corrective submovements) and more end point error.

It is important to note here that not all actions consist of both components. Although each action needs a ballistic component to get the action initiated, actions may be successfully completed in the absence of feedback-based control. There are two common circumstances in which actions are completed without (or with minimal influence from) feedback-based control. The first circumstance in which feedback-based control is not needed is situations in which end point accuracy demands are minimal (e.g., when there is a low index of difficulty, the target is really large and/or close to the effector). Feedback-based corrections might not occur here because the programmed component of the action is accurate enough to achieve the goal. The second circumstance involves situations in which actions are completed in a very short amount of time. Because the feedback loops require time to effectively influence the actions, feedback-based corrections during rapid or ballistic actions are simply not possible. The actor still receives the response-produced information at the end of the movement and can determine whether they have successfully completed the response and can use that information to adjust the next action (i.e., make an offline correction to the action). The information,

however, cannot be used online (during the action) to ensure its accurate completion. Thus, for ballistic actions, such as key presses, a continual source of target information during execution will not affect performance because online corrections cannot be made. Successful completion of action in this context is dependent on the accuracy of the motor program. In contrast, for movements with a longer execution time, such as finger- or mouse-based aiming movements, a continual source of information facilitates accurate completion because the information can be used to make online corrections to the unfolding action.

In sum, the critical implication from this discussion of the use of visual information in motor programming and control is that different types of actions require different types and amounts of information. Specifically, because key press responses are completed in a ballistic manner without the use of feedback, the stable sources of information regarding the target location are not needed to ensure accurate completion. In contrast, because aiming movements generally take longer to complete and have higher accuracy demands, a continual and stable source of visual information about the effector and the target is needed for efficient feedback-based corrections and movement accuracy. As will be discussed later, recent findings suggest that the ways in which we perceive and attend to objects in the world is determined, in part, by the to-be-performed response mode. Thus, careful consideration of response mode is necessary when designing work environment to ensure the efficient extraction of the relevant information and use of the system.

1.1.3 TRANSLATION, CODING, AND MAPPING

As outlined in the preceding sections, the dominant models of human information processing (e.g., Fitts and Posner 1967) distinguishes three basic processes: stimulus identification, response selection, and response programming. While stimulus identification and response programming are functions of stimulus and response properties, respectively, response selection is associated with the translation between stimuli and responses (Welford 1968).

Translation is the seat of the human “interface” between perception and action. Moreover, the effectiveness of translation processes at this interface is influenced to a large extent by the relation between perceptual inputs (e.g., stimuli) and motor outputs (e.g., responses). Since the seminal work of Fitts and colleagues (Fitts and Seeger 1953; Fitts and Deninger 1954), it has been repeatedly demonstrated that errors and choice reaction times to stimuli in a spatial array decrease when the stimuli are mapped onto responses in a spatially “compatible” manner. Fitts and Seeger (1953) referred to this finding as stimulus–response (S–R) compatibility and ascribed it to cognitive codes associated with the spatial locations of elements in the stimulus and response arrays. Presumably, it is the degree of coding and recoding required to map the locations of stimulus and response elements that determine the speed and accuracy of translation and thus response selection (e.g., Wallace 1971).

The relevance of studies of S–R compatibility to the domain of human factors engineering is paramount. It is now well understood that the design of an optimal HCI in which effective S–R translation facilitates fast and accurate responses is largely determined by the manner in which stimulus and response arrays are arranged and mapped onto each other (e.g., Bayerl, Millen, and Lewis 1988; Chapanis and Lindenbaum 1959; Proctor and Van Zandt 1994). As a user, we experience the recalibrating of perceptual-motor space when we take hold of the mouse and move it in a fairly random pattern when we interact with a computer for the first time. Presumably, what we are doing here is attempting to calibrate our actual movements to the resulting virtual movements of the cursor on the screen. Such recalibrations require neural networks and resources that are in addition to those typically activated during direct or standard mapping conditions (Snyder, Batista, and Andersen 1998). Thus, for optimal efficiency of functioning, it seems imperative that the system is designed to require as little recalibration as possible. Again, our contribution to the first edition of this handbook reviews our work on the area of S–R translation and the implications of this work for HCI (Chua, Weeks, and Goodman 2003). We encourage those who are more interested in these issues to read that chapter. For the present chapter, we will instead outline some newer considerations and consequences for contexts in which there is a more direct translation between movements of the user and the effects of these actions in virtual space.

1.2 PERCEPTUAL-MOTOR INTERACTION: ATTENTION AND PERFORMANCE

The vast literature on selective attention and its role in the filtering of target from nontarget information (e.g., Cherry 1953; Treisman 1964a,b, 1986; Deutsch and Deutsch 1963; Treisman and Gelade 1980) has no doubt been informative in the resolution of issues in HCI pertaining to stimulus displays and inputs (e.g., the use of color and sound). However, attention should not be thought of as a unitary function, but rather as a set of information-processing activities that are important for perceptual, cognitive, and motor skills. Indeed, the evolution of HCI into the realm of augmented reality, teleoperation, gestural interfaces, and other areas that highlight the importance of dynamic perceptual-motor interactions, necessitates a greater consideration of the role of attention in the selection and execution of action. Recent developments in the study of how selective attention mediates perception and action and, in turn, how intended actions influences attentional processes, are poised to make just such a contribution to HCI. We will now turn to a review of these developments and some thoughts on their potential relevance to HCI.

1.2.1 ATTENTION

We are all familiar with the concept of “attention” on a phenomenological basis. Even our parents, who likely never formally

studied cognition, demonstrated their understanding of the essential characteristics of attention when they directed us to “pay attention” when we were daydreaming or otherwise not doing what was asked. They knew that humans, like computers, have a limited capacity to process information in that we can only receive, interpret, and act upon a fixed amount of information at any given moment. As such, they knew that any additional, nontask processing would disrupt the performance of our goal-task, be it homework, cleaning, or listening to their lecture. But what is “attention”? What does it mean to “pay attention”? What influences the direction of our attention? The answers to these questions are fundamental to understanding how we interact with our environment. Thus, it is paramount for those who are involved in the design of HCI to consider the characteristics of attention and its interactive relationship with action planning.

1.2.1.1 Characteristics of Attention

Attention is the collection of processes that allow us to dedicate our limited information-processing capacity to the purposeful (cognitive) manipulation of a subset of available information. Stated another way, attention is the process through which information enters into working memory and achieves the level of consciousness. There are three important characteristics of attention: (1) attention is selective and allows only a specific subset of information to enter the limited processing system; (2) the focus of attention can be shifted from one source of information to another; and (3) attention can be divided such that, within certain limitations, one may selectively attend to more than one source of information at a time. The well-known “cocktail party” phenomena (Cherry 1953) effectively demonstrates these characteristics.

Picture yourself at the last busy party or poster session you attended where there was any number of conversations continuing simultaneously. You know from your own experience that you are able to filter out other conversations and *selectively* attend to the single conversation in which you are primarily engaged. You also know that there are times when your attention is drawn to a secondary conversation that is continuing nearby. These *shifts* of attention can occur automatically, especially if you hear your name dropped in the second conversation, or voluntarily, especially when your primary conversation is boring. Finally, you know that you are able to *divide* your attention and follow both conversations simultaneously. However, although you are able to keep track of each discussion simultaneously, you will note that your understanding and contributions to your primary conversation diminish as you dedicate more and more of your attentional resources to the secondary conversation. The diminishing performance in your primary conversation is, of course, an indication that the desired amount of information processing has exceeded your limited capacity.

Although the “cocktail party” example outlined here uses auditory stimuli, the ability to select, divide, and shift attentional resources holds for different modalities (e.g., vision, proprioception) and across multiple modalities (e.g., one can

shift from auditory stimuli to visual stimuli). Because vision is the dominant modality of information transfer in HCI, we will concentrate our discussion on visual selective attention. It should be noted, however, that there is a growing literature on cross-modal influences on attention, especially visual–auditory system interactions (e.g., Spence et al. 2000), that will be relevant in the near future. For those interested in a broader review of the characteristics of attention are encouraged to read our contribution to the second edition of the handbook (Welsh et al. 2007).

1.2.1.2 Shifts of Attention

Structural analyses of the retinal (photosensitive) surface of the eye have revealed two distinct receiving areas—the fovea and the perifoveal (a.k.a. peripheral) areas. The fovea is a relatively small area (about 2° – 3° of visual angle) near the center of the retina, which has the highest concentration of color-sensitive cone cells. It is this high concentration of color-sensitive cells that provides the rich, detailed information that we typically use to identify objects. There are several important consequences of this structural and functional arrangement. First, because of the fovea's pivotal role in object identification and the importance of object identification for the planning of action and many other cognitive processes, visual attention is typically dedicated to the information received by the fovea. Second, because the fovea is such a small portion of the eye, we are unable to derive a detailed representation of the environment from a single fixation. As a result, it is necessary to constantly move information from objects in the environment onto the fovea by rapidly and accurately rotating the eye. These rapid eye movements are known as saccadic eye movements. Because of the tight link between the location of visual attention and saccadic eye movements, these rapid eye movements are referred to as overt shifts of attention.

Although visual attention is typically dedicated to foveal information, it must be remembered that the perifoveal retinal surface also contains color-sensitive cells and, as such, is able to provide details about objects. A *covert* shift of attention refers to any situation in which attention is being dedicated to a nonfoveated area of space. Covert shifts of attention are used when an individual wants or needs to maintain the fovea on a particular object while continuing to scan the remaining environment for other stimuli. Covert shifts of attention also occur immediately before the onset of an overt shift of attention or other type of action (e.g., Shepherd, Findlay, and Hockey 1986). For this reason, people are often able to identify stimuli at the location of covert attention before the acquisition of that location by foveal vision (i.e., overt attention) (Deubel and Schneider 1996).

Both overt and covert shifts of attention can be driven by stimuli in the environment or by the will of the performer. Shifts of attention that are driven by stimuli are known as *exogenous*, or bottom–up, shifts of attention. They are considered to be automatic in nature and thus, for the most part, are outside of cognitive influences. Exogenous shifts of attention are typically caused by a dynamic change in the environment

such as the sudden, abrupt appearance (onset) or disappearance (offset) of a stimulus (e.g., Pratt and McAuliffe 2001), a change in the luminance or color of a stimulus (e.g., Folk, Remington, and Johnston 1992; Posner, Nissen, and Ogden 1978; Posner and Cohen 1984), or the abrupt onset of object motion (e.g., Abrams and Chirst 2003; Folk, Remington, and Wright 1994). The effects of exogenous shifts have a relatively rapid onset, but are fairly specific to the location of the dynamic change and are transient, typically reaching their peak influence around 100 ms after the onset of the stimulus (Cheal and Lyon 1991; Müller and Rabbitt 1989). From an evolutionary perspective, it could be suggested that these automatic shifts of attention developed because such dynamic changes would provide important survival information such as the sudden, unexpected appearance of a predator or prey. However, in more modern times, these types of stimuli can be used to quickly draw one's attention to the location of important information.

In contrast, performer-driven, or *endogenous*, shifts of attention are under complete voluntary control. The effects of endogenous shifts of attention take longer to develop, but can be sustained over a much longer period of time (Cheal and Lyon 1991; Müller and Rabbitt 1989). From an HCI perspective, there are advantages and disadvantages to the fact that shifts of attention can be under cognitive control. The main benefit of cognitive control is that shifts of attention can result from a wider variety of stimuli such as symbolic cues like arrows, numbers, or words. In this way, performers can be cued to locations or objects in the scene with more subtle or permanent information than the dynamic changes that are required for exogenous shifts. The main problem with endogenous shifts of attention is that the act of interpreting the cue requires a portion of the limited information-processing capacity and thus can interfere with, or be interfered by, concurrent cognitive activity (Jonides 1981).

Although it was originally believed that top–down processes could not influence exogenous shifts of attention (i.e., that dynamic changes reflexively capture attention regardless of intention), Folk, Remington, and Johnston (1992) demonstrated that this is not always the case. The task in the Folk et al. (1992) study was to identify a stimulus that was presented in one of four possible locations. For some participants, the target stimulus was a single abrupt onset stimulus (the target appeared in one location and nothing appeared in the other three locations), whereas for the remaining participants the target stimulus was a color singleton (a red stimulus that was presented at the same time as white stimuli that appeared in the other three possible locations). One-hundred and fifty milliseconds before the onset of the target, participants received cue information at one of the possible target locations. The cue information was either abrupt onset stimuli at a single location or color singleton information. Across a series of experiments, Folk et al. (1992) found that the cue tended to increase reaction times to the target stimulus when the cue information was presented at a location that was different from where the target subsequently appeared, indicating that attention had initially been exogenously drawn to

the cue. Importantly, the cue stimuli only interfered with the identification of the target stimulus when the characteristics of cue stimuli matched the characteristics of the target stimulus (i.e., onset cue-onset target and color cue-color target conditions). When the characteristics of the cue did not match the target stimulus (i.e., onset cue-color target and color cue-onset target conditions), the location of the cue did not influence reaction times. Thus, these results reveal that dynamic changes only capture attention when the performer is searching for a dynamic change stimulus. Stated another way, it seems that “automatic” attentional capture is dependent on the expectations of the performer. Folk et al. suggested that people create an attention set in which they establish their expectations for the characteristics of the target stimulus. Stimuli meeting the established set will automatically capture attention, whereas stimuli that do not meet the established set will not.

Subsequent work on this contingent involuntary capture of attention effect has revealed that this attentional set can only be broadly-tuned in that it is most sensitive for discriminating between so-called static (e.g., color singletons) and dynamic (e.g., abrupt onset singletons) discontinuities. For example, Folk, Remington, and Wright (1994) found that a motion singleton (one object suddenly starting to move) and an offset singleton (one object suddenly disappearing) captured attention when participants were searching for an onset singleton target (see also Gibson and Kelsey 1998). Thus, when key press responses are required to a target that is characterized by a dynamic change in the environment, other dynamic change will fit the attentional set and capture attention. The obvious implication of these results is that the most efficient HCIs will be those for which the designer has considered perceptual expectations of the person controlling the system. As we will discuss in Section 1.2.1.3, however, consideration of the perceptual expectations alone is, at best, incomplete.

1.2.1.3 Action-Centered Attention

The majority of the literature reviewed thus far has involved experiments that investigated attentional processes through tasks that used simple or choice key press actions. Cognitive scientists typically use these arbitrary responses because (1) key press responses are relatively uncomplicated and provide simple measures of performance, namely reaction time and error; and (2) by using a simple response, the researcher assumes that they have isolated the perceptual and attentional processes of interest from additional complex motor programming and control processes. Although there are certainly numerous examples of HCI in which the desired response is an individual key press or series of key presses, there are perhaps as many situations in which more complicated movements are required. Indeed, mouse- and joystick-based interactions are in many ways complicated aiming movements. Further, as HCIs move increasingly into virtual reality, touchscreen, tangible interfaces, and other more complex environments, it will become increasingly important to consider the ways in which attention and motor processes interact. Thus, it will become

more critical to determine if the same principles of attention apply when more involved motor responses are required. In addition, some cognitive scientists have suggested that, because human attention systems have developed through evolution to acquire the information required to plan and control complex actions, studying attention under such constrained response conditions may actually provide an incomplete or biased view of attention (Allport 1987, 1993). The tight link between attention and action is apparent when one recognizes that covert shifts of attention occur before saccadic eye movements (Deubel and Schneider 1996) and that overt shifts of attention are tightly coupled to manual aiming movements (Helsen et al. 1998, 2000). Such considerations, in combination with neuro-anatomical studies revealing tight links between the attention and motor centers (Rizzolatti, Riggio, and Sheliga 1994), have led to the development of action-centered models of attention (Rizzolatti et al. 1987; Tipper, Howard, and Houghton 1999; Welsh and Elliott 2004a).

1.2.1.3.1 *The Relationship between Attentional Capture and Action Coding*

Recent research has demonstrated that the behavioral consequences of selecting and executing target-directed actions in the presence of action-relevant nontarget stimuli extend beyond the time taken to prepare and execute the movement (e.g., Meegan and Tipper 1998; Pratt and Abrams 1994). Investigations in our labs and others have revealed that the actual execution of the movement changes in the presence of distractors. For example, there are reports that movements will deviate toward (Welsh, Elliott, and Weeks 1999; Welsh and Elliott 2004a; Welsh et al. 2007; Song and Nakayama 2008; Carr, Phillips, and Meehan 2008; Buetti and Kerzel 2009) or away from (Howard and Tipper 1997; Tipper, Howard, and Jackson 1997; Welsh and Elliott 2004a,b) the nontarget stimulus. For a recent review of the effects of cognitive states on reaching movements, please see Song and Nakayama (2009).

Welsh and Elliott have developed the model of response activation to account for and integrate this research. Consistent with the conclusions of Tipper, Lortie, and Baylis (1992), Welsh and Elliott (2004a) based the model of response activation on the premise that attention and action processes are so tightly linked that the dedication of attention to a particular stimulus automatically initiates response-producing processes that are designed to interact with that stimulus. Responses are activated to attended stimuli regardless of the nature of attentional dedication (i.e., reflexive or voluntary). It is proposed that each time a performer approaches a known scene, a “response set” is established in working memory in which the performer identifies and maintains the characteristics of the expected target stimulus and the characteristics of the expected response to that stimulus. Thus, the response set in the model of response activation is an extension of the attentional set of Folk et al. (1992) in that the response set includes the performer’s expectations of the target stimulus as well as preexcited (preprogrammed) and/or preinhibited

response codes. Each stimulus that matches the physical characteristics established in the response set captures attention and, as a result, activates an independent response process. Stimuli that do not possess at least some of the expected characteristics do not capture attention and thus do not activate responses. Thus, if only one stimulus in the environment matches the response set, then that response process is completed unopposed and the movement emerges rapidly and in an uncontaminated form. However, under conditions in which more than one stimulus matches the response set, multiple response representations are triggered and subsequently race one another to surpass the threshold level of neural activation required to initiate a response. It is important to note that this is not a “winner-take-all” race where only the characteristics of the winning response influence the characteristics of actual movement alone. Instead, the characteristics of the observed movement are determined by the activation level of each of the competing responses at the moment of movement initiation. In this way, if more than one neural representation is active (or if one is active and one is inhibited) at response initiation, then the emerging response will have characteristics of both responses (or characteristics that are opposite to the inhibited response).

The final relevant element of the model is that the activation level of each response is determined by at least three interactive factors—the salience of the stimulus and associated response, an independent inhibitory process, and the time course of each independent process. The first factor, the salience or action-relevancy of the stimulus, is in fact the summation of a number of separate components including the degree attentional capture (based on the similarity between the actual and anticipated stimulus within the response set), the complexity of the response afforded by the stimulus, and the S–R compatibility. When attentional capture and S–R compatibility are maximized and response complexity is minimized, the salience of an individual response is maximized and the response to that stimulus is activated rapidly.

So, what implications does the model of response activation have for the design of HCI? In short, because the model of response activation provides a fairly comprehensive account of movement organization in complex environments, it could be used as the basis for the design of interfaces that consider the cognitive system as an interactive whole as opposed to separate units of attention and movement organization. One of the more obvious implications is that a designer should consider the time intervals between the presentation of each stimulus in a multiple-stimuli set, as this can have dramatic effects on the performer’s ability to quickly respond to each stimulus (e.g., psychological refractory period—Telford 1931; Pashler 1994) and the physical characteristics of each response (Welsh and Elliott 2004a).

1.2.1.3.2 *Spatial Coordinates of Attention in Different Action Contexts*

Arguably the most influential work in the development of the action-centered models was the article by Tipper et al. (1992). Participants in these studies were presented with

nine possible target locations, arranged in a three by three matrix, and were asked to identify the location of a target stimulus appearing at one of these locations while ignoring any nontarget stimuli presented at one of the remaining eight locations. The key innovation of this work was that Tipper and colleagues asked participants to complete a rapid aiming movement to the target location instead of identifying it with a key press. Previous studies of the reference frame of attention using key press responses had revealed that attention can work in retinotopic (e.g., Eriksen and Eriksen 1974), egocentric (e.g., Downing and Pinker 1985; Gawryszewski et al. 1987), and environmental (e.g., Hinton and Parsons 1988) coordinate systems. However, if there is a tight link between attention and action and the requirements of the action modulate, in part, the distribution of attention and attentional capture, then coordinate system used (and subsequent pattern of distractor interference effects observed) during aiming movements should be different from that used during key press responses. This difference in coordinate systems should be observed because the amount and type of information needed to successfully plan and complete aiming movements are different from that needed to successfully complete a key press response (see Section 1.1.2).

Consistent with traditional key press studies, Tipper, Lortie, and Baylis (1992) found that the presence of a distractor increased response times to the target. Although the finding of distractor interference in this selective reaching task was an important contribution to the field in and of itself, the key discovery was that the magnitude of the interference effects caused by a particular distractor location was dependent on the aiming movement being completed. Specifically, it was found that distractors (1) closer to the starting position of the hand (between the start position and the target) cause more interference than distractors farther from the starting position (the proximity-to-hand effect); and, (2) ipsilateral to the moving hand caused more interference than those in the contralateral side of space (the ipsilateral effect). Based on this pattern of interference, Tipper et al. (1992) concluded that attention and action are tightly linked such that the distribution of attention is dependent on the action that was being performed (i.e., attention was distributed in an action-centered coordinate system). Specifically, stimuli that afford actions that are more efficiently executed (i.e., movements of shorter amplitude [Fitts 1954] or into ipsilateral space [Fisk and Goodale 1985]) tend to capture attention to a greater degree (and cause more interference) than distractors that afford less-efficient responses (i.e., movements of longer amplitude or into contralateral space; see also, Tipper, Meegan, and Howard 2002).

Although the study of Tipper et al. (1992) provided critical initial insights into the issue of response efficiency and the action-dependent patterns of interference, additional research has revealed that this pattern of interference is modulated by the characteristics of the environment and the task. For instance, Keulen et al. (2002) have demonstrated that the distance between targets and distractors in the environment alters the attentional frame of reference used during reaching movements. In support of Tipper et al. (1992)

action-centered frame of reference, they found that distractors closer to the start position of the hand caused more interference than distractors beyond the path of the reaching movement (i.e., a proximity-to-hand effect) when there was a large distance (20 mm) between the target and distractor locations. In contrast, when the target and distractor locations were close (5 mm) to each other, a symmetrical pattern of interference was observed in which distractors on either side of the target caused the same amount of interference (i.e., no proximity-to-hand effect was observed). The authors suggested that this shift in the pattern of interference occurred because the planning and control stages of aiming movements require different frames of reference (action-centered and environmental, respectively). These data support the action-centered view in that the patterns interference was even dependent on the stage of action planning and execution. Within the realm of HCI, these data highlight the need for careful consideration of the spatial arrangement of stimuli in the environment because even small changes in the array can alter the efficiency of target engagement.

1.2.1.3.3 *The Capture of Attention in Different Action Contexts*

As reviewed in Section 1.2.1.3.2, initial investigations into action-centered attention were focused primarily on the influence that the spatial location of distractors with respect to the target had on the planning and execution of action (e.g., Meegan and Tipper 1998; Lyons et al. 1999; Pratt and Abrams 1994; Tipper et al. 1992). In that context, an action-centered framework has offered a useful perspective for the spatial organization of perceptual information presented in an HCI context. However, the reason for engaging a target in an HCI task is because the target symbolically represents an outcome or operation to be achieved. Indeed, this is what defines an icon as a target—target features symbolically carry a meaning that defines it as the appropriate target. Whether by intuition and trial and error, or through consideration of the research on attentional capture (e.g., Folk et al. 1992), programmers have already used a variety of dynamic changes to the stimulus characteristics (e.g., suddenly appearing, blinking, moving, growing, etc.) to draw our attention to certain objects and in the hopes of facilitating target engagement. Although there is little doubt that the dynamic stimuli are, in large part, successful in achieving these goals, recent investigations of how the context of the response influence perception and attention suggest that target engagement may be made more efficient through consideration of the response mode, the requirements of the actions system, and the relationship between the stimulus and the desired response.

As an initial illustration of the tight link between perceptual-motor processes, there is a growing body of evidence revealing how the characteristics of the prepared action influence the processing of certain visual stimuli. For instance, Lindemann and Bekkering (2009; see also Craighero et al. 1999) have shown that the degree of

congruency between the action goal and the characteristics of an irrelevant stimulus can facilitate reaction times to initiate the movement. Participants in the study were told to reach out and grasp an X-shaped object as if they were going to turn it clockwise or counterclockwise. They were told in advance which type of movement they would be making and to wait for a “go” signal before initiating the movement that they had prepared. The “go” signal was apparent motion of an object in either a clockwise or counterclockwise direction. It was found that the participants initiated their movements more rapidly when the apparent motion of the “go” signal was congruent with the movement that they had prepared (e.g., a clockwise movement with a clockwise rotating stimulus) than when the apparent motion was incongruent with the prepared movement (e.g., a clockwise movement with a clockwise rotating stimulus). This congruency effect is consistent with other research and demonstrates that prepared movements enhance the perception of characteristics of objects that are related to the to-be-performed movement. For example, the preparation of grasping movements enhances the detection of targets that varied by size, whereas the preparation of pointing movements enhances the detection of targets that varied by luminance (Wykowska, Schubo, and Hommel 2009; see also, Symes et al. 2008).

While the research described in the previous paragraph suggests that perception of specific features is enhanced in an action-specific manner, recent work from our lab suggests that attentional capture is likewise modified by the requirement of the motor system. Specifically, Welsh and Pratt (2008) found that the attentional capture by some dynamic changes is different when key press and aiming responses are required. In this study, participants were asked to identify the location of an onset or offset target stimulus while ignoring a distractor stimulus of the opposite characteristics (i.e., onset targets were paired with offset distractors and vice versa). In separate experiments, participants responded to the target stimulus with a choice key press response or an aiming movement to the target location. Consistent with the findings of Folk et al. (1992) and Folk et al. (1994), interference effects were observed when an offset distractor was presented with an onset target and when an onset distractor was paired with an offset target. When aiming responses were required, however, interference effects were only observed when an onset distractor was presented with an offset target. The offset distractor did not cause an interference effect when participants were aiming to an onset target. Stated another way, the results indicated that an onset distractor slowed responding to an offset target in both key press and aiming tasks. An offset distractor, however, only interfered with task performance when a key press was required.

It was proposed that this action-dependent pattern of interference effects emerged because the action system modified the attentional set, thereby influencing what stimulus features capture attention and those that do not, based on the salience of the stimulus feature for the requirements of the to-be-performed action. Because key press tasks are ballistic in nature, a constant source of stable visual information is

not needed to ensure accurate completion. As a result, any dynamic discontinuity is as salient as any other and can capture attention. In contrast, because the accuracy of aiming movements depends on a continual source of stimulus information for feedback-based control, offset and onset stimuli represent the two extreme ends of saliency to the motor system (with onsets at the maximally salient end and offset at the minimally salient end). As a result, offset stimuli have a very low salience to the motor system and are very unlikely to capture attention when an aiming response is required. In contrast, because onset stimuli are highly salient to the motor system, they are very likely to capture attention when aiming responses are required, regardless of the features of the target stimulus. Thus, it seems that the context of the action and the requirements of the motor system to ensure the accurate completion of the response help to shape the attentional set and what does and does not capture attention (see also Higgins and Welsh, submitted; Welsh and Zbinden 2009).

Similar action-specific interference effects to those observed in our lab have been shown across pointing and grasping actions (Bekkering and Neggers 2002; Weir et al. 2003), pointing and verbal responses (Meegan and Tipper 1999), and different types of pointing responses (Meegan and Tipper 1999; Tipper, Meegan, and Howard 2002). In sum, there is growing evidence that traditional conception of the information-processing stream as serial series of events with action only occurring after perception and cognition stages are completed is in need of revision. The research reviewed here suggests that the action system has what would traditionally be considered as an “upstream” effect and plays an important role in shaping perception and attention. From an applied perspective, now that HCI is moving into virtual reality and other types of assisted response devices, it will become increasingly important to consider the required and/or anticipated action when designing HCI environments. Specifically, this work on the spatial layout (e.g., Keulen et al. 2002) and the characteristics of the stimuli (e.g., Welsh and Pratt 2008) highlights the need for the designer to consider the interactions among perception, attention, and motor processing because there are some situations in which the transfer from simple to complex movements is not always straightforward.

1.2.1.4 Summary

Taken into the realm of HCI, it is our position that the interplay between shifts of attention, spatial compatibility, and object recognition will be a central human performance factor as technological developments continue to enhance the “directness” of direct-manipulation systems (cf. Shneiderman 1983, 1992). Specifically, as interactive environments become better abstractions of reality with greater transparency (Rutkowski 1982), the potential influence of these features of human information processing will likely increase. Thus, it is somewhat ironic that the view toward virtual reality, as the solution to the problem of creating the optimal display representation, may bring with it an “unintended consequence” (Tenner 1996). Indeed, the operator in

such an HCI environment will be subject to the same constraints that are present in everyday life.

The primary goal of human factors research is to guide technological design in order to optimize perceptual-motor interactions between human operators and the systems they use within the constraints of maximizing efficiency and minimizing errors. Thus, the design of machines, tools, interfaces, and other sorts of devices utilizes knowledge about the characteristics, capabilities, as well as limitations, of the human perceptual-motor system. In computing, the development of input devices such as the mouse and graphical user interfaces was intended to improve human–computer interaction. As technology has continued to advance, the relatively simple mouse and graphical displays have begun to give way to exploration of complex gestural interfaces and virtual environments. This development may perhaps, in part, be a desire to move beyond the “artificial” nature of such devices as the mouse, to ones that provide a better mimic of reality. Why move an arrow on a monitor using a hand-held device to point to a displayed object, when instead, you can “reach” and “interact” with the object? Perhaps such an interface would provide a closer reflection of real-world interactions—and the seeming ease with which we interact with our environments, but also subject to the constraints of the human system. With this in mind, we now turn to an alternative approach to perceptual-motor interactions that we believe may point us in some exciting new directions.

1.3 COMMON CODING ACCOUNTS OF PERCEPTUAL-MOTOR INTERACTIONS

At the same time that the research on action-centered attention and perception is gaining momentum, a new approach to cognition has begun to emerge broadly termed “embodied cognition.” This approach argues that, among other things, there is a bidirectional relationship between the body and cognition such that actions are influenced by cognitive operations and cognitive operations are influenced by movements and the body’s action state. In many ways, cognition is considered to be a form of action. One of the key mechanisms that is considered to support this two-way connection between the body and cognition is a common representation in the brain that codes both the action plan and the sensory consequences of the action plan (the effects the action will have on the environment). It is this specific common coding mechanism that differentiates this theory from the modified views of the traditional information-processing theories reviewed in Section 1.2.1.3. On a functional level, it is suggested that these common codes connect the perception, execution, and imagination of movements and, as a result, can also help to shape other cognitive processes. Although there is a literature on the connections between action and a variety of cognitive processes, we will focus here on the relevant literature related to the interactions among action, perception, and imagination.

The origins of the common coding approach can be found in the seminal text of William James (1890). The more modern

and in-depth development of this idea was first articulated by Prinz (1992) and has been refined and expanded as data accumulate (see, Decety 2002; Hommel et al. 2001; Prinz 2005). Simply put, a central outcome of this common coding mechanism is that perception and action are intimately linked such that the activation of one component automatically activates the coupled component. The planning of an action automatically activates a representation of the sensory consequences of the action and, conversely, perception or imagination of an effect automatically activates a representation of the action(s) that can bring about that effect. As a result, one can activate or simulate an action by conceiving of the desired effects on the environment and the effects of a planned action on the environment can be anticipated with the activation or simulation of the response.

A suggested consequence of this coding is that the motor system is activated when humans perceive and imagine movement-related information. This motor system activation and connection between movement (activation of motor representations), observation of movements (activation of perceptual representations), and imagination of movements (covert activation of motor and perceptual representations), then leads to the preferences and biases of our own movements, which can guide the way we perceive and imagine other movements and actions, and may also influence the way we process representations that embed movements (such as verbs). Consistent with these ideas, recent work has extended this effect to language and concept processing, showing that there is motor activation while imagining words encoding movements, and processing sentences involving movements (Bergen, Chang, and Narayan 2004; Wilson and Gibbs 2007; Holt and Beilock 2006; Barsalou 1999).

A common instance of the embodied resonance and simulation process that may involve the common codes is familiar to cinema goers: while watching an actor moving along a precipice, viewers may move their arms and legs or displace body weight to one side or another, based on what they would like to see the actor doing in the scene. Similar effects are seen in sports fans watching athletes perform and novice video game players interacting with their virtual character. Such “simulation” of others’ actions may also underlie our ability to project ourselves into different character roles as well. For instance, this effect may explain why we are emotionally moved by a dramatic film scene: we simulate the characters’ movements and emotional expressions using our own body and, as a result, recreate their emotional states.

In implementation terms, common coding can be thought of as an artificial neural network encoding both action and perception elements, where the activation of one type of element automatically activates the other elements (associative priming), similar to connectionist implementations of semantic priming (Cree, McRae, and McNorgan 1999). Imagination of movement, in this view, would be a form of implicit activation of the action network. Recent modeling work has shown how such common coding could arise purely through agent–environment interactions, when agents move from not using any representations (being purely reactive) to a strategy of

using stored structures in the world/head. In addition, this model shows that common coding can arise from both evolutionary and within-lifetime learning (Chandrasekharan and Stewart 2007).

Most of the evidence for common coding is derived from behavioral studies in which it is assessed how actions in one medium (e.g., imagination) leads to a difference in reaction time or accuracy in another medium (e.g., execution). The following is a brief review of the experimental evidence for different types of interactions. For the sake of space and relevance to HCI, our review will focus on this behavioral evidence for common coding. It should be noted, however, that this behavioral evidence is supported by neurophysiological experiments, including imaging, transcranial magnetic stimulation (TMS), and patient studies (for a comprehensive review, see Rizzolatti and Craighero 2004; or Brass and Heyes 2005). Finally, for the sake of brevity, we will focus our discussion on the less intuitive and more relevant research on the implications of the common coding system for perception–action and imagination–action relationships. We focus on these relationships because we have used this work as the theoretical basis for a collaborative project to develop a novel tangible interface that we will highlight at the end of the chapter.

1.3.1 PERCEPTION–ACTION COMMON CODING

If common coding holds and the perception of an action automatically activates the observed action codes in the observer, then two distinct predictions can be made. The first prediction is that the observation and perception of a movement should negatively influence the concurrent performance of a movement when the observed and executed actions are incompatible because different action codes are activated in the motor system through execution and observation. Thus, the codes of observed action should interfere with the codes of the action that is to be executed. This interference effect would be similar to the trajectory deviation effects caused by competing response codes observed in the action-centered attention studies reviewed earlier in the chapter (e.g., Welsh and Elliott 2004a).

In support of the common coding hypothesis, Kilner, Paulignan, and Blakemore (2003; see also Brass, Bekkering, and Prinz 2001) found that there was more variability in the performance of a rhythmic movement pattern when participants observed another individual performing an incompatible versus a compatible rhythmic pattern. Specifically, when participants were performing a rhythmic up-and-down movement pattern with their arms, there was more horizontal deviation in movement pattern when they observed another person performing a horizontal movement pattern than when the observed person performed a vertical movement pattern. Critically, this interference effect did not occur when the participants observed similar compatible and incompatible movement patterns being executed by a robot arm. This contrast in effects of the human and the robot suggests that the activation of the common codes through observation may be sensitive to the characteristics

of the observed motion and/or the intentionality the observer is able to attribute to the observed actor.

The second, and probably more relevant, prediction is that the perception of actions should be affected by performance of those actions because recent or extensive execution improves the coding of and familiarity with the perceptual consequences of the action. There are a number of lines of evidence that are consistent with this prediction. One line of evidence is the repeated finding that people are better able to recognize actions after having practiced the action patterns. For example, Casile and Giese (2006) found that people were better able to visually recognize a specific movement pattern faster than other movement sequences after learning the movement pattern. Critically, because participants were blindfolded during the learning of the task, the improvement in *visual* recognition was based on verbal and haptic feedback alone. In a related set of studies, Knoblich, Sebanz, and colleagues (see Knoblich and Sebanz 2006 for a review) have shown that people can accurately identify their own action patterns from those of other people. Presumably, people are very accurate at recognizing their own actions because they have a lifetime of experiencing and building of knowledge of their own action–effect relationships.

In addition to the work on recognition, this effect of learned actions seems to extend to preference judgments. When skilled and novice typists were asked to pick between dyads of letters (such as FV and FJ), the skilled typists preferred dyads that would be typed with less interference (i.e., different fingers), whereas novices showed no preference. Moreover, a motor task performed in parallel to the dyad preference judgments lowered skilled typists' preference, but only when the motor task involved the specific fingers that would be used to type the dyads (Beilock and Holt 2007). This preference effect has been generalized recently by Topolinski and Strack (2009), who showed that the mere exposure effect (MEE; stimuli that are repeatedly encountered are increasingly liked) is dependent on motor simulations. They showed that chewing gum while evaluating stimuli destroyed MEEs for words, but not for visual characters. However, kneading a ball with the hand left both MEEs unaffected. They argued that this effect stems from individuals representing stimuli by covertly simulating the sensorimotor processes that run when the stimuli are perceived or acted on. Chewing disrupts this process, kneading does not. These preference effects have recently been used to explain the strong identification players develop with video game characters (Chandrasekharan et al. 2010).

1.3.2 IMAGINATION–ACTION COMMON CODING

We believe the most straightforward and convincing demonstration of the involvement of the motor system in imagination, at least in the imagination of actions, is the repeated finding that the time to mentally execute actions closely corresponds to the time it takes to actually perform them (Decety 2002; Jeannerod 2006; Young, Pratt, and Chau 2009). However, it has also been shown that responses beyond voluntary control

(such as heart and respiratory rates) are activated by imagining actions to an extent proportional to that of actually performing the action (Decety 2002). In sum, these data suggest that imagination of these actions involves the activation of response codes, with these response codes running offline and generating many of the same physiological effects that would be generated during execution, although to a diminished degree. The connections between the motor system and imagination extend beyond the simulation of motor tasks to other cognitive activities (e.g., Hegarty 2004; Martin and Schwartz 2005; Nersessian 2002, 2008). We will center our discussion, however, on mental rotation.

The main prediction of this work is that, if cognitive processes such as imagination and mental rotation engage the common coding system, then these cognitive processes should be affected by concurrent action execution and vice versa. To test the prediction that action planning and execution influences cognition, Wohlschlagel (2001; see also Wexler, Kosslyn, and Berthoz 1998) asked participants to mentally rotate an object while they were planning an action or actually moving their hands or feet in a direction that was compatible or incompatible with the direction of the mental rotation. Consistent with predictions based on the notion of common coding, performance on the mental rotation suffered when the direction of action was incompatible with the direction of mental rotation and performance improved when the direction of action was compatible with the mental rotation.

Although the involvement of our action system in cognition may facilitate efficient processing, the limitations of our motor system may likewise limit or hinder cognitive functioning. For example, it has recently been shown that people with writer's cramp (a focal hand dystonia characterized by constant contractions of the muscles of the hand and forearm that limit hand use) take more time to complete certain mental rotation tasks than their peers without neurological disorders. Interestingly, the difficulties in mental rotation seem to be specific to images of the affected limb (i.e., rotating pictures of hands). The time it took people with focal hand dystonia to rotate pictures of nonbody parts (e.g., houses and cars) were not different from their peers without dystonia (Fiorio, Tinazzi, and Agilotti 2006). Likewise, Kosslyn (1994) reports that participants need more time to perform mental rotations that are physically awkward. These data suggest that common coding may restrict or limit our ability to imagine novel actions and movements. Thus, although our action system may be engaged to facilitate certain cognitive processes, its role is limited by our action repertoire.

1.3.3 SUMMARY

Through this review, we have attempted to concisely summarize the critical features of common coding theories and the evidence that supports these views. Although this area is, in many ways, in its infancy, there is a clear growing body of evidence supporting a common code system linking execution, perception, and imagination of movement and

that this system can be accessed to support a wide variety of cognitive processes. The vast majority of the research in this area has been directed to testing and expanding the theoretical aspects of common coding. We believe, however, that there is tremendous potential for the principles outlined in common coding theory to shape and enhance HCI. In fact, this theoretical approach has recently been used to derive novel embodied interaction designs. We will describe this development and some potential applications in the second half of the following section.

1.4 PERCEPTUAL-MOTOR INTERACTION IN APPLIED TASKS: A FEW EXAMPLES

As we mentioned at the outset of this chapter, the evolution of computers and computer-related technology has brought us to the point at which the manner with which we interact with such systems has become a research area in itself. Current research in motor behavior and experimental psychology pertaining to attention, perception, action, and spatial cognition is poised to make significant contributions to the area of HCI. In addition to the continued development of a knowledge base of fundamental information pertaining to the perceptual-motor capabilities of the human user, these contributions will include new theoretical and analytical frameworks that can guide the study of HCI in various settings. In this final section, we highlight just a few specific examples of HCI situations that offer a potential arena for the application of the basic research that we have outlined in this chapter.

1.4.1 ATTENTION CUEING FOR MILITARY TARGET DETECTION

Combat identification of friends and enemies is essential for mission effectiveness and the prevention of friendly fire. The software that projects images to the operator's displays, including images from unmanned aerial vehicles and on head mounted displays (HMDs), can cue attention to possible target locations. In each situation, the user is required to navigate and engage targets in the real world, while attempting to perform a detection and identification on the interactive display. The use of this assistive software creates a dual task in which the operator must divide his or her attention between the separate tasks in order to complete the job successfully. Although the identification cues can provide great opportunities to facilitate the detection of critical information, they could also decrease performance by creating distracting clutter on the display (Yeh and Wickens 2001a,b). These effects are magnified as the cue reliability decreases and, in the case where the task can be performed easily without the cue, it has been shown that imperfect cues may hinder performance (Maltz and Shinar 2003). In this context, an error of commission (i.e., the cue indicates a nontarget) has much greater behavioral consequences than an error of omissions (i.e., the technology fails to cue a target) (Maltz and Shinar 2003).

Thus, it is imperative that the cue stimuli involved in the secondary identification task be carefully designed to ensure the efficient processing of this information to allow as much of the attentional resources as possible to be available for the real-world tasks of target engagement.

An HMD can assist with target detection because it overlays critical cue information over the actual environment, reducing the scanning time required to sample and attend both the display and the environment. An HMD also allows for cueing in x and y coordinates and the use of *conformal imagery* in which cues or information is presented in a world-referenced frame rather than a screen-referenced frame eliminating the need for the user to transfer between reference frames (Yeh, Wickens and Seagull 1999). Users are also better at recovering from cueing errors using an HMD (Yeh et al. 2003). However, HMDs are especially susceptible to the detrimental effects of clutter as the user is expected to attend concurrently to information both on the display and in the environment.

The majority of studies find that cueing assists the user in detecting the target more quickly and accurately (e.g., Maltz and Shinar 2003); however, there is often a cost for detecting uncued targets (e.g., Yeh et al. 2003) and an increase in false alarms (Yeh and Wickens 2001a,b). The cost may result from "attentional tunnelling" where the participants fail to direct their attention to areas outside of the cue. The tunnelling may result from the user creating an attention set (Folk et al. 1992; Folk et al. 1994) for specific cue features. This attentional set may increase the chances of these salient cue stimuli capturing attention, but at the same time reduce the chances that stimuli not in the set (i.e., uncued targets) capture attention. Overall, cueing can assist the user in directing attention in difficult detection tasks as long as the cue is sufficiently reliable and does not induce clutter into the visual scene.

1.4.2 DERIVING NOVEL INTERACTION DESIGNS FROM COMMON CODING

Although there are clear implications for the research outlined above for the design of stimuli in virtual environments, we also believe that the principles of perceptual-cognitive-motor interactions outlined above should shape and enhance the interface devices that are used to translate our action goals into virtual environments. In fact, two of us (Timothy N. Welsh and Sanjay Chandrasekharan) have been involved in the recent development of a novel interaction device (see Mazalek et al. 2010). The goal of the research was to develop a device that more effectively mapped the actions of the user to the movements of the avatar. The rationale for this goal being that, by translating the user's own actions onto the avatar, there will be a shorter recalibration period, and the user can more easily relate to the avatar and respond more efficiently in the virtual environment. An additional, yet to be tested, potential consequence of this more direct relationship between user and avatar, is that once the user has identified with the avatar's movements, it is possible that the user can

then learn from the avatar if it moves in a novel action pattern that is physically consistent with the movements of the user. Thus, a more direct translation from user to avatar might not only facilitate performance in virtual environments, but it might also assist in the learning and development of the user.

The development of this device is rooted in the common coding theory (e.g., Prinz 1997). A particularly informative set of findings are that individuals recognize their own actions more accurately than the action patterns of other people, even when all that is available is a very information-poor rendition such as point-light displays (see Knoblich and Sebanz 2006, for a review). This own-action advantage is thought to arise because the observer's motor system is involved in the perception process. Because the motor system of the observer is trained to their own actions and the sensory consequences of those actions, it is thought that viewing their own actions more efficiently activates their motor system and the tightly linked common perceptual codes. This more efficient activation of the common codes then allows the individuals to identify their own actions more accurately than those of other people. Extending these findings to HCI, we reasoned that a user would identify more closely with a virtual character in a virtual environment if that virtual character encodes the player's own actions as opposed to movement primitives common to all, or at least a subset of, characters.

Based on this experimental and theoretical work, a control interface was developed to more directly map the user's own actions onto a virtual character in a real-time virtual environment. The device that was developed was a wearable, jointed puppet whose limbs are attached to the limbs of the user so that the limbs of the puppet move along with the hands, legs, and neck of the user. Potentiometers are located at the joints so that the changes in joint angle of the puppet can be transferred to a virtual character (Mazalek et al. 2010). As an initial testing of the puppet system, we recently examined if the same "own-action" advantage (people can recognize their own movements better than others, Knoblich and Sebanz 2006) was present when their movements were represented by a virtual character. Consistent with previous work, we have found that individuals were able to identify abstract representations (point-light displays) of their own actions when the representations were created by affixing small lights to actor's actual body (see Mazalek et al. 2009) and when a player's movements are transferred to an avatar using the puppet (Mazalek et al. 2010). The advantages persist even when the point-light walkers were presented in altered body sizes (Mazalek et al. 2009). Thus, we feel confident that the movements of an individual can be effectively transferred to virtual characters through the puppet device and that people may be able to identify with (embody) these characters when this transfer of movement patterns is successful.

Although our initial development and testing of the device seems positive, the interface continues to evolve. As the interface improves, our view to the possible applications of this system, beyond real-time interaction, expands. For example, we are opening a second line of research in which we are trying to exploit the link between action,

cognition, and imagination. Extending the results from the research reviewed above and the theoretical relationship between action, cognition, and imagination, we hypothesized that novel movements executed by the embodied avatar may improve imagination of novel movements, thus improving players' ability to execute creative cognitive processes such as mental rotation. To facilitate this learning effect, however, the "embodied" virtual characters (characters encoding the player's own actions) would need to execute movements on screen that are impossible for the actual user to perform. Further, the user will lose some control of the embodied avatar when the avatar executes novel movements, such as back-flips (as this would require the user also doing back-flips). Thus, the puppet-controlled avatar will need to retain the movement patterns of the user while executing these physically impossible movements. This is an interesting application challenge, where we need to maintain a fine line between control and no-control, with self-recognition elements of the former situation retained/continued into the latter situation.

As an initial attempt to solve this issue, we have developed a game in which the cameras around the avatar rotate slowly, giving the impression of the avatar rotating in space. Objects then appear close to the avatar, and the user's task is to touch these objects using the puppet interface. Our preliminary results reveal that playing this game using the puppet leads to improved performance on the game and a mental rotation task compared with playing the game using standard game interfaces, such as keyboards and game controllers. These and other experimental applications are still under development. While we are hopeful that the puppet device will achieve all our aims, we feel that, regardless of the outcome, this entire line of research is a powerful example of how theoretical considerations of perceptual-cognitive-motor interactions can be used to inform HCI development and, likewise, this technological development can lead to new methods for testing and enhancing the theory on which the technology was based. For a wider discussion of how common coding theory can help in deriving novel interaction modes, see Chandrasekharan et al. (2010).

1.5 SUMMARY

The field of HCI offers a rich environment for the study of perceptual-motor interactions. The design of effective human–computer interfaces has been, and continues to be, a significant challenge that demands an appreciation of the entire human perceptual-motor system. The information-processing approach has provided a dominant theoretical and empirical framework for the study of perceptual-motor behavior in general, and for consideration of issues in HCI and human factors in particular. Texts in the area of human factors and HCI (including the present volume) are united in their inclusion of chapters or sections that pertain to the topic of human information processing. Moreover, the design of effective interfaces reflects our knowledge of the perceptual (e.g., visual displays, use of sound, graphics), cognitive (e.g., conceptual models, desktop metaphors), and motoric

constraints (e.g., physical design of input devices, ergonomic keyboards) of the human perceptual-motor system.

Technological advances have undoubtedly served to improve the HCI experience. For example, we have progressed beyond the use of computer punch cards and command-line interfaces to more complex tools such as graphical user interfaces, speech recognition, and tangible control systems. As HCI has become not only more effective, but by the same token more elaborate, the importance of the interaction between the various perceptual, cognitive, and motor constraints of the human system has come to the forefront. In our previous chapters, we presented overviews of some topics of research in action-centered attention and in S-R compatibility in perceptual-motor interactions that we believed were relevant to HCI. In the present chapter, we have added an overview of common coding theories of cognition. We believe that the relevance of the research and theoretical considerations discussed in this chapter for HCI cannot be underestimated. Clearly, considerable research will be necessary to evaluate the applicability of both of these potentially relevant lines of investigation to specific HCI design problems. Nevertheless, the experimental work to date leads us to conclude that the motor system is not simply responsible for outputting the results of perceptual and cognitive processing, but in fact has a critical and active role in shaping perception and cognition. For this reason, an effective interface must be sensitive to the perceptual and action expectations of the user, the specific action associated with a particular response location, the action relationship between that response and those around it, and the degree of translation required to map the perceptual-motor workspaces.

ACKNOWLEDGMENTS

We would like to recognize the financial support of the Natural Sciences and Engineering Research Council of Canada, Alberta Ingenuity Fund, and the National Science Foundation.

REFERENCES

- Abrams, R. A., and S. E. Chirst. 2003. Motion onset captures attention. *Psychol Sci* 14:427–32.
- Adam, J. J., R. Mol, J. Pratt, and M. H. Fischer. 2006. Moving farther but faster: An exception to Fitts's law. *Psychol Sci* 17:794–8.
- Adams, J. A. 1971. A closed-loop theory of motor learning. *J Mot Behav* 3:111–50.
- Adams, J. A., and Dijkstra, S. 1966. Short-term memory for motor responses. *J Exp Psychol* 71:314–8.
- Allport, A. 1987. Selection for action: Some behavioural and neurophysiological considerations of attention and action. In *Perspectives on Perception and Action*, ed. H. Heuer, and A. F. Sanders. 395–419. Hillsdale, NJ: Erlbaum.
- Allport, A. 1993. Attention and control: Have we been asking the wrong questions? A critical review of twenty-five years. In *Attention and Performance 14: Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience*, ed. D. E. Meyer and S. Kornblum, 183–218. Cambridge, MA: MIT Press.
- Barsalou, L. W. 1999. Perceptual symbol systems. *Behav Brain Sci* 22:577–660.
- Bayerl, J., D. Millen, and S. Lewis. 1988. Consistent layout of function keys and screen labels speeds user responses. In *Proceedings of the Human Factors Society 32nd Annual Meeting*, 334–46. Santa Monica, CA: Human Factors Society.
- Beilock, S. L., and L. E. Holt. 2007. Embodied preference judgments: Can likeability be driven by the motor system? *Psychol Sci* 18:51–7.
- Bekkering, H., and S. F. W. Neggers. 2002. Visual search is modulated by action intentions. *Psychol Sci* 13:370–4.
- Bergen, B., N. Chang, and S. Narayan. 2004. Simulated action in an embodied construction grammar. In *Proceedings of the 26th Annual Meeting of the Cognitive Science Society*, ed. K. D. Forbus, D. Gentner, and T. Regier, 108–13. Hillsdale, NJ: Lawrence Erlbaum.
- Brass, M., H. Bekkering, and W. Prinz. 2002. Movement observation affects movement execution in a simple response task. *Acta Psychol* 106:3–22.
- Brass, M., and C. Heyes. 2005. Imitation: Is cognitive neuroscience solving the correspondence problem? *Trends Cogn Sci* 9:489–95.
- Broadbent, D. E. 1958. *Perception and Communication*. New York, NY: Pergamon.
- Brooks, V. B. 1974. Some examples of programmed limb movements. *Brain Res* 71:299–308.
- Buetti, S., and D. Kerzel. 2009. Conflicts during response selection affect response programming: Reactions toward the source of stimulation. *J Exp Psychol Hum Percept Perform* 35:816–34.
- Card, S. K., W. K. English, and B. J. Burr. 1978. Evaluation of mouse, rate-controlled isometric joystick, step keys, and text keys for text selection on a CRT. *Ergonomics* 21:601–13.
- Carr, S. M., J. G. Phillips, and J. W. Meehan. 2008. Non-target flanker effects on movement in a virtual action centered reference frame. *Exp Brain Res* 184:95–103.
- Casile, A., and M. A. Giese. 2006. Non-visual motor learning influences the recognition of biological motion. *Curr Biol* 16:69–74.
- Chandrasekharan, S., A. Mazalek, Y. Chen, M. Nitsche, and A. Ranjan. 2010. Ideomotor Design: using common coding theory to derive novel video game interactions. *Pragmat Cogn* 18:313–39.
- Chandrasekharan, S., and T. C. Stewart. 2007. The origin of epistemic structures and proto-representations. *Adapt Behav* 15:329–53.
- Chapanis, A., and L. E. Lindenbaum. 1959. A reaction time study of four control-display linkages. *Hum Factors* 1:1–7.
- Cheal, M. L., and D. R. Lyon. 1991. Central and peripheral precuing of forced-choice discrimination. *Q J Exp Psychol* 43A:859–80.
- Cherry, E. C. 1953. Some experiments on the recognition of speech, with one and with two ears. *J Acoust Soc Am* 25:975–9.
- Chua, R., and D. Elliott. 1993. Visual regulation of manual aiming. *Hum Mov Sci* 12:365–401.
- Chua, R., D. J. Weeks, and D. Goodman. 2003. Perceptual-Motor Interaction: Some Implications for HCI. In *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies, and Emerging Applications*, ed. J. A. Jacko, and A. Sears, 23–34. Hillsdale, NJ: Lawrence Earlbaum Assoc.
- Craighero, L., L. Fadiga, G. Rizzolatti, and C. Umiltà. 1999. Action for perception: A motor-visual attentional effect. *J Exp Psychol Hum Percept Perform* 25:1673–92.
- Cree, G. S., K. McRae, and C. McNorgan. 1999. An attractor model of lexical conceptual processing: Simulating semantic priming. *Cogn Sci* 23:371–414.

- Crossman, E. R. F. W., and P. J. Goodeve. 1983. Feedback control of hand movement and Fitts' Law. *Paper presented at the meeting of the Experimental Psychology Society*, Oxford, July 1963. Published in *Q J Exp Psychol* 35A:251–78.
- Decety, J. 2002. Is there such a thing as a functional equivalence between imagined, observed and executed actions. In *The Imitative Mind: Development, Evolution and Brain Bases*, ed. A. N. Meltzoff and W. Prinz, 291–310. Cambridge: Cambridge University Press.
- Deubel, H., and W. X. Schneider. 1996. Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Vis Res* 36:1827–37.
- Deutsch, J. A., and D. Deutsch. 1963. Attention: Some theoretical considerations. *Psychol Rev* 70:80–90.
- Downing, C. J. and S. Pinker. 1985. The spatial structure of visual attention. In *Attention and Performance XI*, ed. M. I. Posner and O. S. M. Marin, 171–87. Hillsdale, NJ: Erlbaum.
- Elliott, D., and M. A. Khan. 2010. *Vision and Goal-directed Movements: Neurobehavioral Perspectives*. Champaign, IL: Human Kinetics.
- Elliott, D., R. G. Carson, D. Goodman, and R. Chua. 1991. Discrete vs. continuous visual control of manual aiming. *Hum Mov Sci* 10:393–418.
- English, W. K., D. C. Engelhart, and M. L. Berman. 1967. Display-selection techniques for text manipulation. *IEEE Trans Hum Factors Electron* 8:5–15.
- Eriksen, B. A., and C. W. Eriksen. 1974. Effects of noise letters upon the identification of a target letter in a non-search task. *Percept Psychophys* 16:143–6.
- Fiorio, M., M. Tinazzi, and S. M. Agilotti. 2006. Selective impairment of hand mental rotation in patients with focal hand dystonia. *Brain* 129:47–54.
- Fisk, J. D., and M. A. Goodale. 1985. The organization of eye and limb movements during unrestricted reaching to targets in contralateral and ipsilateral space. *Exp Brain Res* 60:159–78.
- Fitts, P. M. 1954. The information capacity of the human motor system in controlling the amplitude of movement. *J Exp Psychol* 47:381–91.
- Fitts, P. M. 1964. Perceptual-motor skills learning. In *Categories of Human Learning*, ed. A. W. Melton, 243–85. New York: Academic Press.
- Fitts, P. M., and C. M. Seeger. 1953. S-R compatibility: Spatial characteristics of stimulus and response codes. *J Exp Psychol* 46:199–210.
- Fitts, P. M., and M. I. Posner. 1967. *Human Performance*. Belmont, CA: Brooks-Cole.
- Fitts, P. M., and R. I. Deninger. 1954. S-R compatibility: Correspondence among paired elements within stimulus and response codes. *J Exp Psychol* 48:483–91.
- Folk, C. L., R. W. Remington, and J. C. Johnson. 1992. Involuntary covert orienting is contingent on attentional control settings. *J Exp Psychol Hum Percept Perform* 18:1030–44.
- Folk, C. L., R. W. Remington, and J. H. Wright. 1994. The structure of attentional control: Contingent attentional capture by apparent motion, abrupt onset, and color. *J Exp Psychol Hum Percept Perform* 20:317–29.
- Gawryszewski, L., L. Riggio, G. Rizzolatti, and C. Umiltá. 1987. Movements of attention in the three spatial dimensions and the meaning of “neutral” cues. *Neuropsychologia* 25:19–29.
- Gibson, B. S., and E. M. Kelsey. 1998. Stimulus-driven attentional capture is contingent on attentional set for display-wide visual features. *J Exp Psychol Hum Percept Perform* 24:699–706.
- Heath, M. 2005. Role of limb and target vision in the online control of memory-guided reaches. *Motor Control* 9:281–311.
- Hegarty, M. 2004. Mechanical reasoning as mental simulation. *Trends Cogn Sci* 8:280–5.
- Helsen, W. F., D. Elliott, J. L. Starkes, and K. L. Ricker. 1998. Temporal and spatial coupling of point of gaze and hand movements in aiming. *J Mot Behav* 30:249–59.
- Helsen, W. F., D. Elliott, J. L. Starkes, and K. L. Ricker. 2000. Coupling of eye, finger, elbow, and shoulder movements during manual aiming. *J Mot Behav* 32:241–8.
- Hick, W. E. 1952. On the rate of gain of information. *Q J Exp Psychol* 4:11–26.
- Hinton, G. E., and L. A. Parsons. 1988. Scene-based and viewer-centered representations for comparing shapes. *Cognition* 30:1–35.
- Hoffman, E. R., and J. T. A. Lim. 1997. Concurrent manual-decision tasks. *Ergonomics* 40:293–318.
- Holt, L. E., and S. L. Beilock. 2006. Expertise and its embodiment: Examining the impact of sensorimotor skill expertise on the representation of action-related text. *Psychon Bull Rev* 13:694–701.
- Hommel, B., J. Müsseler, G. Aschersleben, and W. Prinz. 2001. The theory of event coding (TEC): A framework for perception and action planning. *Behav Brain Sci* 24:849–78.
- Hornecker, E., R. J. K. Jacob, C. Hummels, B. Ullmer, A. Schmidt, E. V. D. Hoven, and A. Mazalek. 2008. TEI goes on: Tangible and embedded interaction. *IEEE Pervasive Comput* 7:91–6.
- Howard, L. A., and S. P. Tipper. 1997. Hand deviations away from visual cues: Indirect evidence for inhibition. *Exp Brain Res* 113:144–52.
- Hyman, R. 1953. Stimulus information as a determinant of reaction time. *J Exp Psychol* 45:188–96.
- James, W. 1890. *Principles of Psychology*. New York: Holt.
- Jeannerod, M. 2006. From volition to agency: The mechanism of action recognition and its failures. In *Disorders of Volition*, ed. N. Sebanz and W. Prinz, 175–92. Cambridge, MA: The MIT Press.
- Jonides, J. 1981. Voluntary versus automatic control over the mind's eye's movement. In *Attention and Performance IX*, ed. J. Long and A. Baddley, 187–203. Hillsdale, NJ: Lawrence Erlbaum Assoc.
- Kelso, J. A. S. 1982. The process approach to understanding human motor behavior: An introduction. In *Human Motor Behavior: An Introduction*, ed. J. A. S. Kelso, 3–19. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Kelso, J. A. S. 1995. *Dynamic Patterns: The Self-organization of Brain and Behavior*. Cambridge, MA: MIT Press.
- Kelso, J. A. S., D. L. Southard, and D. Goodman. 1979. On the coordination of two-handed movements. *J Exp Psychol Hum Percept Perform* 5:229–38.
- Keulen, R. F., J. J. Adam, M. H. Fischer, H. Kuipers, and J. Jolles. 2002. Selective reaching: Evidence for multiple frames of reference. *J Exp Psychol Hum Percept Perform* 28:515–26.
- Kilner, J. M., Paulignan, Y., Blakemore, S-J. 2003. An interference effect of observed biological movement on action. *Current Biology* 13:522–525.
- Khan, M. A., I. M. Franks, D. Elliott, G. P. Lawrence, R. Chua, P. M. Bernier, S. Hansen, and D. Weeks. 2006. Inferring online and offline processing of visual feedback in target-directed movements from kinematic data. *Neurosci Biobehav Rev* 30:1106–21.
- Knoblich, G., and N. Sebanz. 2006. The social nature of perception and action. *Psychol Sci* 15:99–104.
- Kosslyn, S. M. 1994. *Image and Brain*. Cambridge, MA: The MIT Press.

- Lindemann, O., and H. Bekkering. 2009. Object manipulation and motion perception: Evidence of an influence of action planning on visual processing. *J Exp Psychol Hum Percept Perform* 35:1062–71.
- Lyons, J., D. Elliott, K. L. Ricker, D. J. Weeks, and R. Chua. 1999. Action-centred attention in virtual environments. *Can J Exp Psychol* 53:176–78.
- MacKenzie, C. L., R. G. Marteniuk, C. Dugas, D. Liske, and B. Eickmeier. 1987. Three dimensional movement trajectories in Fitts' task: Implications for control. *Q J Exp Psychol* 39A:629–47.
- Maltz, M., and D. Shinar. 2003. New alternative methods of analyzing human behavior in cued target acquisition. *Hum Factors* 45:281–95.
- Marteniuk, R. G. 1976. *Information Processing in Motor Skills*. New York: Holt, Rinehart and Winston.
- Marteniuk, R. G., C. L. MacKenzie, M. Jeannerod, S. Athenes, and C. Dugas. 1987. Constraints on human arm movement trajectories. *Can J Psychol* 41:365–78.
- Marteniuk, R. G., C. L. MacKenzie, and J. L. Leavitt. 1988. Representational and physical accounts of motor control and learning: Can they account for the data? In *Cognition and Action in Skilled Behavior*, ed. A. M. Colley and J. R. Beech, 173–90. Amsterdam: Elsevier Science Publishers.
- Martin, T., and D. L. Schwartz. 2005. Physically distributed learning: Adapting and reinterpreting physical environments in the development of fraction concepts. *Cogn Sci* 29:587–625.
- Mazalek, A., M. Nitsche, S. Chandrasekharan, T. Welsh, P. Clifton, A. Quitmeyer, F. Peer, and F. Kirschner. 2010. Recognizing self in puppet controlled virtual avatars. In *Proceedings of Fun and Games 2010*, 66–73. New York: ACM Press.
- Mazalek, A., S. Chandrasekharan, M. Nitsche, T. Welsh, G. Thomas, T. Sanka, and P. Clifton. 2009. Giving your self to the game: Transferring a player's own movements to avatars using tangible interfaces. In *Proceedings of Sandbox 2009: ACM SIGGRAPH Videogame Symposium*, 161–8. New York, NY: ACM Press.
- Meegan, D. V., and S. P. Tipper. 1998. Reaching into cluttered visual environments: Spatial and temporal influences of distracting objects. *Q J Exp Psychol* 51A:225–49.
- Meegan, D. V., and S. P. Tipper. 1999. Visual search and target-directed action. *J Exp Psychol Hum Percept Perform* 25:1347–62.
- Meyer, D. E., R. A. Abrams, S. Kornblum, C. E. Wright, and J. E. K. Smith. 1988. Optimality in human motor performance: Ideal control of rapid aimed movements. *Psychol Rev* 95:340–70.
- Müller, H. J., and P. M. A. Rabbitt. 1989. Reflexive and voluntary orienting of visual attention: Time course of activation and resistance to interruption. *J Exp Psychol Hum Percept Perform* 15:315–30.
- Nersessian, N. J. 2002. The cognitive basis of model-based reasoning in science. In *The Cognitive Basis of Science*, ed. P. Carruthers, S. Stich, and M. Siegal, 133–53. Cambridge: Cambridge University Press.
- Nersessian, N. J. 2008. *Creating Scientific Concepts*. Cambridge, MA: The MIT Press.
- Pashler, H. 1994. Dual-task interference in simple tasks: Data and theory. *Psychol Bull* 116:220–44.
- Posner, M. I. 1982. Cumulative development of attentional theory. *Am Psychol* 37:168–79.
- Posner, M. I., and Y. Cohen. 1984. Components of visual orienting. In *Attention and Performance X*, ed. H. Bouma and D. G. Bouwhuis, 531–56. Hillsdale, NJ: Lawrence Earlbaum Assoc.
- Posner, M. I., and S. W. Keele. 1969. Attentional demands of movement. In *Proceedings of the 16th Congress of Applied Psychology*. Amsterdam: Swets and Zeitlinger.
- Posner, M. I., M. J. Nissen, and W. C. Ogden. 1978. Attended and unattended processing modes: The role of set for spatial location. In *Modes of Perceiving and Processing Information*, ed. H. Pick and E. Saltzman, 137–57. Hillsdale, NJ: Lawrence Earlbaum Assoc.
- Pratt, J., and J. McAuliffe. 2001. The effects of onsets and offsets on visual attention. *Psychol Res* 65:185–91.
- Pratt, J., and R. A. Abrams. 1994. Action-centered inhibition: Effects of distractors in movement planning and execution. *Hum Mov Sci* 13:245–54.
- Prinz, W. 1992. Why don't we perceive our brain states? *Eur J Cogn Psychol* 4:1–20.
- Prinz, W. 1997. Perception and action planning. *Eur J Cogn Psychol* 9:129–54.
- Prinz, W. 2005. A common coding approach to imitation. In *Perspectives on Imitation: From Neuroscience to Social Science*, Vol. 1, ed. S. Hurley and N. Chater, 141–56. Cambridge, MA: The MIT Press.
- Proctor, R. W., and T. Van Zandt. 1994. *Human Factors in Simple and Complex Systems*. Boston: Allyn and Bacon.
- Proctor, R. W., and K. L. Vu. 2003. Human information processing. In *The Human-computer Interaction Handbook: Fundamentals, Evolving Technologies, and Emerging Applications*, ed. J. A. Jacko and A. Sears, 35–50. Hillsdale, NJ: Lawrence Earlbaum Assoc.
- Rizzolatti, G., and L. Craighero. 2004. The mirror-neuron system. *Annual Reviews in Neuroscience* 27:169–92.
- Rizzolatti, G., L. Riggio, J. Dascola, and C. Umiltà. 1987. Reorienting attention across the horizontal and vertical meridians: Evidence in favor of a premotor theory of attention. *Neuropsychologia* 25:31–40.
- Rizzolatti, G., L. Riggio, and B. M. Sheliga. 1994. Space and selective attention. In *Attention and Performance XV*, ed. C. Umiltà and M. Moscovitch, 231–65. Cambridge: MIT Press.
- Rutkowski, C. 1982. An introduction to the human applications standard computer interface, part I: Theory and principles. *Byte* 7:291–310.
- Shepherd, M., J. M. Findlay, and R. J. Hockey. 1986. The relationship between eye-movements and spatial attention. *Q J Exp Psychol* 38A:475–91.
- Shneiderman, B. 1983. Direct manipulation: A step beyond programming languages. *IEEE Comput* 16:57–69.
- Shneiderman, B. 1992. *Designing the User Interface: Strategies for Effective Human-computer Interaction*. Reading, MA: Addison-Wesley Publishing Company.
- Snyder, L. H., A. P. Batista, and R. A. Andersen. 1998. Change in motor plan, without a change in the spatial locus of attention, modulates activity in posterior parietal cortex. *J Neurophysiol* 79:2814–19.
- Song, J. H., and K. Nakayama. 2008. Target selection in visual search as revealed by movement trajectories. *Vis Res* 48:853–61.
- Song, J. H., and K. Nakayama. 2009. Hidden cognitive states revealed in choice reaching tasks. *Trends Cogn Sci* 13:360–6.
- Spence, C., D. Lloyd, F. McGlone, M. E. R. Nichols, and J. Driver. 2000. Inhibition of return is supramodal: A demonstration between all possible pairings of vision, touch, and audition. *Exp Brain Res* 134:42–8.
- Stelmach, G. E. 1982. Information-processing framework for understanding human motor behavior. In *Human Motor Behavior: An Introduction*, ed. J. A. S. Kelso, 63–91. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Symes, E., M. Tucker, R. Ellis, L. Vainio, and G. Ottoboni. 2008. Grasp preparation improves change detection for congruent objects. *J Exp Psychol Hum Percept Perform* 34:854–71.

- Telford, C. W. 1931. The refractory phase of voluntary and associative responses. *J Exp Psychol* 14:1–36.
- Tenner, E. 1996. *Why Things Bite Back: Technology and the Revenge of Unintended Consequences*. New York: Alfred A. Knopf.
- Tipper, S. P., C. Lortie, and G. C. Baylis. 1992. Selective reaching evidence for action-centered attention. *J Exp Psychol Hum Percept Perform* 18:891–905.
- Tipper, S. P., D. Meegan, and L. A. Howard. 2002. Action-centred negative priming: Evidence for reactive inhibition. *Vis Cogn* 9:591–614.
- Tipper, S. P., L. A. Howard, and G. Houghton. 1999. Behavioral consequences of selection form neural population codes. In *Attention and Performance XVIII*, ed. S. Monsell and J. Driver, 223–45. Cambridge, MA: MIT Press.
- Tipper, S. P., L. A. Howard, and S. R. Jackson. 1997. Selective reaching to grasp: Evidence for distractor interference effects. *Vis Cogn* 4:1–38.
- Topolinski, S., and F. Strack. 2009. Motormouth: Mere exposure depends on stimulus-specific motor simulations. *J Exp Psychol Learn Mem Cogn* 35:423–33.
- Treisman, A. M. 1964a. The effect of irrelevant material on the efficiency of selective listening. *Am J Psychol* 77:533–46.
- Treisman, A. M. 1964b. Verbal cues, language, and meaning in selective attention. *Am J Psychol* 77:206–19.
- Treisman, A. M. 1986. Features and objects in visual processing. *Scientific Am* 255:114–25.
- Treisman, A. M., and G. Gelade. 1980. A feature-integration theory of attention. *Cogn Psychol* 12:97–136.
- Wallace, R. J. 1971. S-R compatibility and the idea of a response code. *J Exp Psychol* 88:354–60.
- Weir, P. L., D. J. Weeks, T. N. Welsh, D. Elliott, R. Chua, E. A. Roy, and J. Lyons. 2003. Action-centred distractor effects in discrete control selection. *Exp Brain Res* 149:207–13.
- Welford, A. T. 1968. *Fundamentals of Skill*. London: Methuen.
- Welsh, T. N., and D. Elliott. 2004a. Movement trajectories in the presence of a distracting stimulus: Evidence for a response activation model of selective reaching. *Q J Exp Psychol* 57(A):1031–57.
- Welsh, T. N., and D. Elliott. 2004b. The effects of response priming and inhibition on movement planning and execution. *J Mot Behav* 36:200–11.
- Welsh, T. N., D. Elliott, J. G. Anson, V. Dhillon, D. J. Weeks, J. L. Lyons, and R. Chua. 2005. Does Joe influence Fred's action? Inhibition of return across different nervous systems. *Neurosci Lett* 385:99–104.
- Welsh, T. N., D. Elliott, and D. J. Weeks. 1999. Hand deviations towards distractors: Evidence for response competition. *Exp Brain Res* 127:207–12.
- Welsh, T. N., and J. Pratt. 2008. Actions modulate attentional capture. *Q J Exp Psychol* 61:968–76.
- Welsh, T. N., D. J. Weeks, R. Chua, and D. Goodman. 2007. Perceptual-motor interaction: Some implications for HCI. In *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*, 2nd ed., ed. A. Sears and J. A. Jacko, 27–41. Boca Raton, FL: CRC Press.
- Welsh, T. N. and M. Zbinden. 2009. Fitts' Law and action-centered reference frames in selective reaching: The "proximity-to-hand" effect revisited. *Motor Control* 13:100–12.
- Wexler, M., S. M. Kosslyn, and A. Berthoz. 1998. Motor processes in mental rotation. *Cognition* 68:77–94.
- Wilson, N. L., and R. W. Gibbs. 2007. Real and imagined body movement primes metaphor comprehension. *Cogn Sci* 31:721–31.
- Wohlschlagel, A. 2001. Mental object rotation and the planning of hand movements. *Percept Psychophys* 63:709–18.
- Woodworth, R. S. 1899. The accuracy of voluntary movements. *Psychol Rev* 3(Monograph Suppl):1–119.
- Wykowska, A., A. Schubö, and B. Hommel. 2009. How you move is what you see: Action planning biases selection in visual search. *J Exp Psychol Hum Percept Perform* 35:1755–69.
- Yeh, M., and C. D. Wickens. 2001a. Attentional filtering in the design of electronic map displays: A comparison of color coding, intensity coding, and decluttering techniques. *Hum Factors* 43:543–62.
- Yeh, M., and C. D. Wickens. 2001b. Display signalling in augmented reality: Effects of cue reliability and image realism on attention allocation and trust calibration. *Hum Factors* 43:355–65.
- Yeh, M., C. D. Wickens, and F. J. Seagull. 1999. Target cuing in visual search: The effects of conformality and display location on the allocation of visual attention. *Hum Factors* 41:524–42.
- Yeh, M., J. L. Merlo, C. D. Wickens, and D. L. Bradenburg. 2003. Head up versus head down: The costs of imprecision, unreliability, and visual clutter on cue effectiveness for display signalling. *Hum Factors* 45:390–407.
- Young, S. J., J. Pratt, and T. Chau. 2009. Misperceiving the speed-accuracy tradeoff: Imagined movements and perceptual decisions. *Exp Brain Res* 192:121–32.

2 Human Information Processing

An Overview for Human–Computer Interaction

Robert W. Proctor and Kim-Phuong L. Vu

CONTENTS

2.1	Human Information-Processing Approach	22
2.2	Information-Processing Methods	22
2.2.1	Signal Detection Methods and Theory	22
2.2.2	Chronometric Methods	23
2.2.3	Speed–Accuracy Methods	24
2.2.4	Psychophysiological and Neuroimaging Methods	25
2.3	Information-Processing Models	26
2.4	Information Processing in Choice-Reaction Tasks	26
2.4.1	Stimulus Identification	27
2.4.2	Response Selection	28
2.4.3	Response Execution	29
2.5	Memory in Information Processing	30
2.5.1	Short-Term (Working) Memory	30
2.5.2	Long-Term Memory	31
2.5.3	Other Factors Affecting Retrieval of Earlier Events	32
2.6	Attention in Information Processing	33
2.6.1	Models of Attention	33
2.6.2	Automaticity and Practice	35
2.7	Problem Solving and Decision Making	35
2.8	Summary and Conclusion	36
	References	37

It is natural for an applied psychology of human-computer interaction to be based theoretically on information-processing psychology.

—Card, Moran, and Newell (1983)

Human–computer interaction (HCI) is fundamentally an information-processing task. When interacting with a computer or any technological device, a user has specific goals and subgoals in his or her mind. For example, smartphone users initiate the interaction by turning on or activating the device and selecting the appropriate commands needed to accomplish their desired goal. Given that the smartphone can do more than just make calls, the commands may activate applications designed to allow specific types of tasks such as playing games, e-mailing, navigating with GPS, or web surfing to be performed. The resulting output, typically displayed on the phone’s screen, must provide adequate

information for the user to complete the next step, or the user must enter another command to obtain the desired output. The sequence of interactions to accomplish the goals may be long and complex, and several alternative sequences, differing in efficiency, may be used to achieve these goals. During the interaction, the user is required to identify displayed information, select responses based on the displayed information, and execute those responses by entering commands. The user must search the displayed information and attend to the appropriate aspects of it. She or he must also recall the commands and the resulting consequences of those commands for different programs, remember information specific to the task that is being performed, and make decisions and solve problems during the process. For the interaction between the device and user to be efficient, the interface must be designed in accordance with the user’s information-processing capabilities.

2.1 HUMAN INFORMATION-PROCESSING APPROACH

The rise of the human information-processing approach in psychology is closely coupled with the growth of the fields of cognitive psychology, human factors, and human engineering (see Proctor and Vu 2010). Although research that can be classified as falling within these fields has been conducted since the last half of the nineteenth century, their formalization dates back to World War II (see Hoffman and Defenbacher 1992). As part of the war efforts, experimental psychologists worked along with engineers on applications associated with using the sophisticated equipment being developed. As a consequence, the psychologists were exposed not only to applied problems but also to the techniques and views being developed in areas such as communications engineering (see Roscoe 2011). Many of the concepts from engineering, for instance, the notion of transmission of information through a limited capacity communications channel, were seen as applicable to analyses of human performance.

The human information-processing approach is based on the idea that human performance, from displayed information to response, is a function of several processing stages. The nature of these stages, how they are arranged, and the factors that influence how quickly and accurately a particular stage operates, can be discovered through appropriate research methods. It is often said that the central metaphor of the information-processing approach is that a human is like a computer (e.g., Lachman, Lachman, and Butterfield 1979). However, even more fundamental than the computer metaphor is the assumption that the human is a complex system that can be analyzed in terms of subsystems and their interrelation. This point is evident in the work of researchers on attention and performance, such as Paul Fitts (1951) and Donald Broadbent (1958), who were among the first to adopt the information-processing approach in the 1950s.

The systems perspective underlies not only human information processing but also human factors and HCI, providing a direct link between the basic and applied fields (Proctor and Van Zandt 2008). Human factors, in general, and HCI in particular, begin with the fundamental assumption that a human–machine system can be decomposed into machine and human subsystems, each of which can be analyzed further. The human information-processing approach provides the concepts, methods, and theories for analyzing the processes involved in the human subsystem. Posner (1986) stated, “Indeed, much of the impetus for the development of this kind of empirical study stem from the desire to integrate description of the human with overall systems” (p. V-6). Young, Clegg, and Smith (2004) emphasized that the most basic distinction among three processing stages (perception, cognition, and action), as captured in a block diagram model of human information processing, is important even for understanding the dynamic interactions of an operator with a vehicle for purposes of computer-aided augmented cognition.

They note,

“This block diagram model of the human is important because it not only models the flow of information and commands between the vehicle and the human, it also enables access to the internal state of the human at various parts of the process. This allows the modeling of what a cognitive measurement system might have access to (internal to the human), and how that measurement might then be used as part of a closed-loop human-machine interface system” (pp. 261–262).

In the first half of the twentieth century, the behaviorist approach predominated in psychology, particularly in the United States. Within this approach, many sophisticated theories of learning and behavior were developed that differed in various details (Bower and Hilgard 1981). However, the research and theories of the behaviorist approach tended to minimize the role of cognitive processes and were of limited value to the applied problems encountered in World War II. The information-processing approach was adopted because it provided a way to examine topics of basic and applied concern such as attention that were relatively neglected during the behaviorist period. It continues to be the main approach in psychology, although contributions have been made from other approaches.

Within HCI, human information-processing analyses are used in two ways. First, empirical studies evaluate the information-processing requirements of various tasks in which humans use computers. Second, computational models are developed with the intent to characterize human information processing when interacting with computers and to predict human performance with alternative interfaces. In this chapter, we survey methods used to study human information processing and summarize the major findings and the theoretical frameworks developed to explain them. We also tie the methods, findings, and theories to HCI issues to illustrate their use.

2.2 INFORMATION-PROCESSING METHODS

Any theoretical approach makes certain presuppositions and tends to favor some methods and techniques over others. Information-processing researchers have used behavioral and, to an ever-increasing extent, psychophysiological and neuroimaging measures, with an emphasis on chronometric (time-based) methods. There also has been a reliance on flow models that are often quantified through computer simulation or mathematical modeling.

2.2.1 SIGNAL DETECTION METHODS AND THEORY

One of the most useful methods for studying human information processing is that of signal detection (Macmillan and Creelman 2005). In a signal detection task, some event is classified as a signal, and the subject’s task is to detect whether the signal is present. Trials on which it is not present are called noise trials. The proportion of trials on which the signal is correctly identified as present is called the hit rate, and the proportion of trials on which the signal is incorrectly

identified as present is called the false alarm rate. By using the hit and false alarm rates, whether the effect of a variable on detectability or response bias can be evaluated.

Signal detection theory is often used as the basis for analyzing data from such tasks. This theory assumes that the response on each trial is a function of two discrete operations, encoding and decision. On a trial, the subject samples the information presented and decides whether this information is sufficient to warrant a *signal present* response. The sample of information is assumed to provide a value along a continuum of evidence states regarding the likelihood that the signal was present. The noise trials form a probability distribution of states, as do the signal trials. The decision that must be made on each trial can be characterized as whether the event is from the signal or noise distribution. The subject is presumed to adopt a criterion value of evidence above which he or she responds *signal present* and below which he or she responds *signal absent*.

In the simplest form, the distributions are assumed to be normal and equal variance. In this case, a measure of detectability, d' , can be derived, as well as a measure of response bias, C (for criterion; Macmillan and Creelman 2005). The d' measure represents the difference in the means for the signal and noise distributions in standard deviation units and is found by converting the hit rate and false alarm rate to standard normal scores and obtaining the difference. A value of 0 indicates no detectability, whereas a value of 3.0 or greater indicates close to perfect detectability. The C measure is calculated by summing the standardized values of the hit and false alarm rates, and dividing by two. A value of 0 for C indicates no response bias. Positive values indicate a bias toward *signal absent* responses, and negative values indicate a bias toward *signal present* responses, with the absolute value indicating the magnitude of the bias. This measure reflects the observer's overall willingness to say *signal present*, regardless of whether it actually is present. There are numerous alternative measures of detectability and bias based on different assumptions and theories, and many task variations to which they can be applied (see Macmillan and Creelman 2005).

Signal detection analyses have been particularly useful because they can be applied to any task that can be depicted in terms of binary discriminations. For example, the proportion of words in a memory task correctly classified as old can be treated as a hit rate, and the proportion of new lures classified as old can be treated as a false alarm rate (e.g., Rotello and Macmillan 2006). In cases such as these, the resulting analysis helps researchers determine whether variables are affecting detectability of an item as old or response bias.

An area of research in which signal detection methods have been widely used is that of vigilance (Parasuraman and Davies 1977). In a typical vigilance task, a display is monitored for certain changes in it (e.g., the occurrence of an infrequent stimulus). Vigilance tasks are common in the military, but many aspects also can be found in computer-related tasks such as monitoring computer network operations (Percival and Noonan 1987). A customary finding for vigilance tasks is the vigilance decrement, in which the hit rate decreases as time on the task increases. The classic example of this vigilance

decrement is that, during World War II, British radar observers detected fewer of the enemy's radar signals after 30 minutes in a radar observation shift (Mackworth 1948). Parasuraman and Davies concluded that, for many situations, the primary cause of the vigilance decrement is an increasingly strict response criterion. That is, the false alarm rate as well as the hit rate decreases as a function of time on task.

Parasuraman and Davies (1977) also provided evidence that detectability decreases across the vigil when the task requires comparison of each event to a standard held in memory and the event rate is high. Findings indicate that this decrease in detectability is a consequence of the high demand on cognitive resources imposed by such tasks. Although vigilance tasks were previously thought to be undemanding, evidence has shown that maintaining a vigil in many situations requires considerable mental effort (Warm, Parasuraman, and Matthews 2008). Our point here is that signal detection theory has played a prominent role in this research on vigilance, helping to dissociate changes in performance associated with mental demands (decreased detectability) from those due to lapses of attention (response criteria).

2.2.2 CHRONOMETRIC METHODS

Chronometric methods, for which time is a factor, have been the most widely used for studying human information processing. Indeed, Lachman, Lachman, and Butterfield (1979) portrayed reaction time (RT) as the main dependent measure of the information-processing approach. Although many other measures are used, RT still predominates in part because of its sensitivity and in part because of the sophisticated techniques that have been developed for analyzing RT data.

A technique called the subtractive method, introduced by Donders (1868/1969) in the 1860s, was revived in the 1950s and 1960s. This method provides a way to estimate the duration of a particular processing stage. The assumption of the subtractive method is that a series of discrete processing stages intervene between stimulus presentation and response execution. Through selection of pairs of tasks that differ by a single stage, the RT for the easier task can be subtracted from that for the more difficult task to yield the time for the additional process. Donders used three tasks hypothesized to differ with respect to stimulus identification and response selection, respectively, and estimated the time for each stage. Recently, Van de Laar et al. (2010) applied similar logic to situations in which on some trials a participant receives a "stop" signal during the reaction process, indicating that the response is to be stopped. They estimated the durations of the stop-signal identification process and a response-mapping process to be 34 and 20 ms, respectively.

The subtractive method has been used to estimate the durations of a variety of other processes, including rates of mental rotation (approximately 12–20 ms per degree of rotation; Shepard and Metzler 1971) and memory search (approximately 40 ms per item; Sternberg 1969). An application of the subtractive method to HCI would be, for example, to compare the time to find a target link on two web pages

that are identical except for the number of links displayed, and to attribute the extra time to the additional visual search required for the more complex web page.

The subtractive method is only applicable when discrete, serial processing stages can be assumed. Also, the processing for the two tasks being compared must be the same except for the additional process that differentiates them. This requires an assumption of pure insertion, which is that the additional process for the more complex of two tasks can be inserted without affecting the processes held in common by the tasks. However, this assumption often is not justified.

Sternberg (1969) developed the additive factors method to allow determination of the processes involved in performing a task. The additive factors method avoids the problem of pure insertion because the crucial data are whether two variables affect RT for the same task in an additive or interactive manner. Sternberg assumed, as did Donders, that information processing occurs in a sequence of discrete stages, each of which produces a constant output that serves as input to the next stage in the sequence. With these assumptions, he showed that two variables that affect different stages should have additive effects on RT. In contrast, two variables that affect the same stage should have interactive effects on RT. Sternberg performed detailed analyses of memory search tasks in which a person holds a set of letters or digits in memory and responds to a target stimulus by indicating whether it is in the memory set. Based on the patterns of additive and interactive effects that he observed, Sternberg concluded that the processing in such tasks involves four stages: target identification, memory search, response selection, and response execution. Grobelny, Karwowski, and Drury (2005) provide an application of additive factors logic to usability of graphical icons in the design of HCI interfaces. Mode of icon array (menu or dialog box), number of icons, and difficulty of movement had additive effects on response times, implying that these variables affect different processing stages.

Both the subtractive and additive factors methods have been challenged on several grounds (Pachella 1974). First, the assumption of discrete serial stages with constant output is difficult to justify in many situations. Second, both methods rely on analyses of RT, without consideration of error rates. This can be problematic because performance is typically not error free, and, as described in Section 2.2.3, speed can be traded for accuracy. Despite these limitations, the methods have proved to be robust and useful (Sanders 1998). For example, Salthouse (2005) notes that the process analysis approach used in contemporary research into aging effects on cognitive abilities “has used a variety of analytical methods such as subtraction, additive factors ... to partition the variance in the target variable into theoretically distinct processes” (p. 288).

2.2.3 SPEED–ACCURACY METHODS

The function relating response speed to accuracy is called the speed–accuracy trade-off (Pachella 1974). The function, illustrated in Figure 2.1, shows that very fast responses




FIGURE 2.1 Speed-accuracy operating characteristic curve. Faster responding occurs at the cost of lower accuracy.

can be performed with chance accuracy, and accuracy will increase as responding slows down. Of importance is the fact that when accuracy is high, as in most RT studies, a small increase in errors can result in a large decrease in RT. With respect to text entry on computing devices, MacKenzie and Soukoreff (2002) state, “Clearly, both speed and accuracy must be measured and analyzed.... Participants can enter text more quickly if they are willing to sacrifice accuracy” (pp. 159–160).

In speed–accuracy trade-off studies, the speed–accuracy criterion is varied between blocks of trials or among subjects by using different instructions regarding the relative importance of speed versus accuracy, varying payoffs such that speed or accuracy is weighted more heavily, or imposing different response deadlines (Wickelgren 1977). These studies have the potential to be more informative than RT studies because they can provide information about whether variables affect the intercept (time at which accuracy exceeds chance), asymptote (the maximal accuracy), and rate of ascension from the intercept to the asymptote, each of which may reflect different processes. For example, Boldini, Russo, and Avons (2004) obtained evidence favoring dual-process models of recognition memory over single-process models by varying the delay between a visually presented test word and a signal to respond. Recognition accuracy benefited from a modality match at study and test (better performance when the study words were also visual rather than auditory) at short response-signal delays, but it benefited from deep processing during study (judging pleasantness) over shallow processing (repeating aloud each word) at long response-signal delays. Boldini et al. interpreted these results as consistent with the view that recognition judgments are based on a fast familiarity process or a slower recollection process.

In tasks requiring search of complex visual displays, a speed emphasis may influence more than just the criterion for emitting a response. McCarley (2009) had young adults perform a simulated baggage-screening task under instructions that emphasized speed or accuracy of responding. With speed emphasis, the participants made fewer eye fixations of shorter duration than under accuracy emphasis. Reduction in accuracy was a consequence mainly of failure

to fixate the target of the search rather than a failure to respond to targets that were fixated. This study illustrates how a speed–accuracy trade-off manipulation can be of value in applied contexts.

Because the speed–accuracy criterion is manipulated in addition to any other variables of interest, much more data must be collected in a speed–accuracy study than in a typical RT study. Consequently, use of speed–accuracy methods has been restricted to situations in which the speed–accuracy relation is of major concern or of apparent significant value, rather than being widely adopted as the method of choice.

2.2.4 PSYCHOPHYSIOLOGICAL AND NEUROIMAGING METHODS

In the past decade, psychophysiological and neuroimaging methods have been used increasingly to evaluate implications of information-processing models and to relate the models to brain processes. This area of research is called *cognitive neuroscience* (Ward 2010). Such methods can provide details regarding the nature of processing by examining physiological activity as a task is being performed. The most widely used psychophysiological method involves measurement of electroencephalograms (EEGs), which are recordings of changes in brain activity as a function of time as measured from electrodes placed on the scalp (Rugg and Coles 1995). Different frequency bands of EEG rhythms can be distinguished that can be related to subjective states and the processes underlying task performance.

One application of EEGs to HCI in recent years has been the development of brain–computer interfaces that allow a person to control technological devices through the use of brain signals. Such interfaces are of value for motor-disabled persons who are not able to communicate through traditional data-entry devices. Changes in EEGs that arise from different types of mental processing can be coded into distinct computer commands, and people can be trained to use their thoughts to control the computer's interface (e.g., Kauhanen et al. 2007). This mode of HCI opens up possibilities for disabled persons to interact with their environment and communicate with other people.

Of most concern for information-processing research are event-related potentials (ERPs), which are the changes in brain activity that are elicited by an event such as stimulus presentation or response initiation. ERPs are obtained by averaging across many trials of a task to remove background EEG noise and are thought to reflect postsynaptic potentials in the brain. There are several features of the ERP that represent different aspects of processing. These features are labeled according to their polarity, positive (P) or negative (N), and their sequence or latency. The first positive (P1) and negative (N1) components are associated with early perceptual processes. They are called exogenous components because they occur in close temporal proximity to the stimulus event and have a stable latency with respect to it. Later components reflect cognitive processes and are called endogenous

because they are a function of the task demands and have a more variable latency than the exogenous components. One such component that has been studied extensively is the P3 (or, P300), which represents postperceptual processes. When an occasional target stimulus is interspersed in a stream of standards, the P3 is observed in response to targets, but not to standards. By comparing the effects of task manipulations on various ERP components such as P3, their onset latencies, and their scalp distributions, relatively detailed inferences about the cognitive processes can be made.

An early application of P3 analysis to HCI is a study by Trimmel and Huber (1998). In their study, subjects performed three HCI tasks (text editing, programming, and playing the game *Tetris*) for 7 minutes each. They also performed comparable paper/pencil tasks in three other conditions. The P3 was measured after each experimental task by having subjects monitor a stream of high- and low-pitched tones, keeping count of each separately. The P3 varied as a function of type of task, as well as medium (computer vs. paper/pencil). The amplitude of the P3 was smaller following the HCI tasks than following the paper/pencil tasks, suggesting that the HCI tasks caused more fatigue or depletion of cognitive resources than the paper/pencil task. The P3 latency was shorter after the programming task than after the others, which the authors interpreted as an aftereffect of highly focused attention.

Another measure that has been used in studies of human information processing is the lateralized readiness potential (LRP; Eimer 1998). The LRP can be recorded in choice-reaction tasks that require a response with the left or right hand. It is a measure of differential activation of the lateral motor areas of the visual cortex that occurs shortly before and during execution of a response. The asymmetric activation favors the motor area contralateral to the hand making the response, because this is the area that controls the hand. The LRP has been obtained in situations in which no overt response is ever executed, allowing it to be used as an index of covert, partial response activation. The LRP is thus a measure of the difference in activity from the two sides of the brain that can be used as an indicator of covert reaction tendencies, to determine whether a response has been prepared even when it is not actually executed. It can also be used to determine whether the effects of a variable are before or subsequent to response preparation.

Electrophysiological measurements do not have the spatial resolution needed to provide precise information about the brain structures that produce the recorded activity, although advances in the technology are producing continual improvements in this regard. Much work has been done recently, though, on neuroimaging methods that provide better spatial resolution. These include positron-emission tomography, functional magnetic resonance imaging (fMRI), and transcranial Doppler sonography, which measure changes in blood flow associated with neuronal activity in different regions of the brain (Huettel, Song, and McCarthy 2004). Traditionally, these methods have poorer temporal resolution

than the electrophysiological methods, but with the introduction of more sophisticated techniques, the gap in temporal resolution has been greatly reduced.

In an imaging study, often both control and experimental tasks are performed, and the functional neuroanatomy of the cognitive processes is derived by subtracting the image during the control task from that during the experimental task. This subtractive method of neuroimaging analysis has the same limitations as that for reaction-time analysis (Sartori and Umiltà 2000). Stevenson, Kim, and James (2009) provided evidence that an additive factors analysis of fMRI data, in which interactive vs. additive effects of different independent variables are compared, “provides a method for investigating multisensory interactions that goes beyond what can be achieved with more established metric-based, subtraction-type methods” (p. 183).

Application of cognitive neuroscience to human factors and HCI has been advocated under the heading of neuroergonomics (e.g., Lees et al. 2010). According to Parasuraman (2003), “Neuroergonomics focuses on investigations of the neural bases of mental functions and physical performance in relation to technology, work, leisure, transportation, health care and other settings in the real world” (p. 5). Neuroergonomics has the goal of using knowledge of the relation between brain function and human performance to design interfaces and computerized systems that are sensitive to brain function with the intent of increasing the efficiency and safety of human–machine systems.

2.3 INFORMATION-PROCESSING MODELS

It is common to assume that the processing between stimuli and responses consists of a series of discrete stages for which the output for one stage serves as the input for the next, as Donders and Sternberg assumed. This assumption is made for the Model Human Processor (Card, Moran, and Newell 1983) and the Executive-Process Interactive Control (EPIC; Meyer and Kieras 1997) architectures, among others, both of which have been applied to HCI. However, models can be developed that allow for successive processing stages to operate concurrently. McClelland’s (1979) cascade model, in which partial information at one subprocess, or stage, is transferred to the next, is of this type. Each stage is continuously active, and its output is a continuous value that is always available to the next stage. The final stage results in selection of which of the possible alternative responses to execute. Many parallel distributed processing, or neural network, models are of a continuous nature.

According to J. Miller (1988), models of human information processing can be classified as discrete or continuous along three dimensions: representation, transformation, and transmission. Representation refers to whether the input and output codes for the processing stage are continuous or discrete. Transformation refers to whether the operation performed by the processing stage (e.g., spatial transformation) is continuous or discrete. Transmission is classified as discrete if the processing of successive stages does not overlap temporally. The discrete stage model proposed by Sternberg

(1969) has discrete representation and transmission, whereas the cascade model proposed by McClelland (1979) has continuous representation, transmission, and transformation. Models can be intermediate to these two extremes. For example, Miller’s (1988) asynchronous discrete coding model assumes that most stimuli are composed of features, and these features are identified separately. Discrete processing occurs for feature identification, but once a feature is identified, this information can be passed to response selection while the other features are still being identified.

Sequential sampling models are able to account for both RT and accuracy, and consequently, the trade-off between them (Ratcliff and Smith 2004; Van Zandt, Colonius, and Proctor 2000). Such models are dynamic models of signal detection, in which decisions are based on a series of samples from the probability distributions rather than a single sample. Each sample is classified as favoring one alternative or another, and this information is fed into a decision mechanism in which gradual accumulation of the information occurs until a response threshold is reached, at which time that response is made. As Busemeyer and Diederich (2010) note, “Dynamic models of signal detection have proven to be very effective for simultaneously analyzing choice probability and the response time distributions for signal detection tasks” (p. 89).

Various types of the dynamic models have been developed and applied to an array of experimental tasks. In such models, factors that influence the quality of information processing (the detectability or discriminability) have their effects on the rate at which the information accumulates. In contrast, factors that bias speed versus accuracy or factors that produce biases toward particular responses have their effects on the response thresholds.

Sequential sampling can be incorporated into more complete cognitive architectures to model speed and accuracy. These architectures specify properties of various processing stages and stores, such as memory and decision processes, and provide a means for developing specific models to simulate performance of a range of tasks. One widely used architecture of this type is Adaptive Control of Thought—Rational (Anderson et al. 2004), which has been used to model, for example, improvements in performance and retention with practice (practice and retention [Anderson, Fincham, and Douglass 1999] and the choices in decision-making tasks [Gonzalez, Lerch, and Lebiere 2003]).

2.4 INFORMATION PROCESSING IN CHOICE-REACTION TASKS

In a typical choice-reaction task in which each stimulus is assigned to a unique response, it is customary to distinguish between three stages of processing: stimulus identification, response selection, and response execution (Proctor and Van Zandt 2008). The stimulus-identification stage involves processes that are entirely dependent on stimulus properties. The response-selection stage concerns those processes involved in determining which response to make to each stimulus. Response execution refers to programming and execution




FIGURE 2.2 Information-processing stages and variables that affect them, based on Sanders' (1998) taxonomy. (From Sanders, A. F., *Elements of Human Performance*, Erlbaum, Mahwah, New Jersey, 1998. With permission.)

of motor responses. Based on additive factors logic, Sanders (1998) decomposed the stimulus-identification stage into three subcategories and the response-execution stage into two subcategories, resulting in six stages (see Figure 2.2).

2.4.1 STIMULUS IDENTIFICATION

The preprocessing stage of stimulus identification refers to peripheral sensory processes involved in the conduction of the sensory signal along the afferent pathways to the sensory projection areas of the cerebral cortex. These processes are affected by variables such as stimulus contrast and retinal location. As stimulus contrast, or intensity, decreases, RT increases. For example, Miles and Proctor (2009) had participants make left and right keypress responses to the nonspatial or spatial feature of centrally presented location words. The discriminability of the spatial feature of the word, or of

both the spatial and nonspatial features, was manipulated. When the spatial feature of the word was task-irrelevant, decreasing the discriminability of this feature reduced the typical benefit for correspondence of the word meaning with the key press response to the relevant stimulus feature. This correspondence benefit was restored when the discriminability of both the task-relevant and task-irrelevant features were reduced together, slowing the processing of both the relevant and irrelevant information. These results suggest that reduction of discriminability slows processing of the perceptual information but does not alter the response-selection processes that operate on that information.

Feature extraction involves lower-level perceptual processing based in area V1 (the visual cortex) and other early visual cortical areas. Stimulus discriminability, word priming, and stimulus quality affect the feature extraction process. For example, manipulations of stimulus quality such as

superimposing a grid slow RT, presumably by creating difficulty for the extraction of features. Identification itself is influenced by word frequency and mental rotation. The latter refers to that when a stimulus is rotated from the upright position, the time it takes to identify the stimulus increases as an approximately linear function of angular deviation from upright (Shepard and Metzler 1971; see also Section 2.2.2). This increase in identification time is presumed to reflect a normalization process by which the image is mentally rotated in a continuous manner to the upright position.

2.4.2 RESPONSE SELECTION

Response selection refers to those processes involved in determining what response to make to a particular stimulus. It is affected by the number of alternatives, stimulus–response compatibility, and precuing (providing advance information about a forthcoming event). RT increases as a logarithmic function of the number of stimulus–response alternatives (Hick 1952; Hyman 1953). This relation is known as the Hick–Hyman law, which for N equally likely alternatives is as follows:

$$RT = a + b \log_2 N \quad (2.1)$$

where a is the base processing time, and b is the amount that RT increases with increases in N . The slope of the Hick–Hyman function is influenced by many factors. For example, the slope decreases as subjects become practiced at a task (Teichner and Krebs 1974). Usher, Olami, and McLelland (2002) provided evidence from fits of a sequential sampling model that the Hick–Hyman law is due to subjects’ adjusting their response criteria upward as the number of alternatives increases, in an attempt to maintain a constant high level of accuracy.

One variable that influences the slope of the Hick–Hyman function is stimulus–response compatibility, which has considerable impact on response-selection efficiency (see Proctor and Vu 2006, for a review of compatibility principles). Compatibility effects are differences in speed and accuracy of responding as a function of how natural, or compatible, the relation between stimuli and responses is. Two types of compatibility effects can be distinguished (Kornblum, Hasbroucq, and Osman 1990). For one type, certain sets of stimuli are more compatible with certain sets of responses than with others. For example, the combinations of verbal–vocal and spatial–manual sets yield better performance than the combinations of verbal–manual and spatial–vocal sets (Wang and Proctor 1996). For the other type, within a specific stimulus–response set, some mappings of individual stimuli to responses produce better performance than others. If one stimulus has the meaning “left” and the other “right,” performance is better if the left stimulus is mapped to the left response and the right stimulus to the right response, for all stimulus and response modes.

Fitts and Seeger (1953) and Fitts and Deininger (1954) demonstrated both types of compatibility effects for spatially

arranged display and response panels. However, compatibility effects occur for a much wider variety of other stimulus–response sets. According to Kornblum, Hasbroucq, and Osman (1990), dimensional overlap (similarity) between the stimulus and response sets is the critical factor. When the sets have dimensional overlap, a stimulus will activate its corresponding response automatically. If this response is correct (compatible mapping), responding will be facilitated, but if it is not correct (incompatible mapping), responding will be inhibited. A second factor contributing to the advantage for the compatible mapping is that intentional translation of the stimulus into a response will occur quicker when the mapping is compatible than when it is not. Most contemporary models of stimulus–response compatibility include both automatic and intentional response-selection routes (Hommel and Prinz 1997), although they differ regarding the exact conditions under which each plays a role and the way in which they interact.

One reason why automatic activation is considered to contribute to compatibility effects is that such effects occur when irrelevant stimulus information overlaps with the response set (Lu and Proctor 1995). The Stroop color-naming effect, for which an incongruent color word produces interference in naming a relevant stimulus color, is most well-known example. An irrelevant stimulus location also produces interference when it is incongruent with the location of a key press to a relevant stimulus dimension, a phenomenon known as the Simon effect (Simon 1990). Psychophysiological studies in which the LRP has been measured have provided evidence that the Simon effect is due, at least in part, to activation of the response corresponding to stimulus location (Melara et al. 2008).

For completely unrelated stimulus and response sets that are structured, performance is better when structural correspondence is maintained (Reeve and Proctor 1990). For instance, when stimuli and responses are ordered (e.g., a row of four stimulus locations and a row of four response locations), RT is faster when the stimulus–response mapping can be characterized by a rule (e.g., press the key at the mirror opposite location) than when the mapping is random (Duncan 1977). Spatial compatibility effects also occur when display and response elements refer to orthogonal spatial dimensions (Proctor and Cho 2006). However, stimulus–response compatibility effects sometimes do not occur under conditions in which one would expect them to. For example, when compatible and incompatible mappings are mixed within a single block, the typical compatibility effect is eliminated (Shaffer 1965; Vu and Proctor 2004). Moreover, the same display and response elements can be coded along multiple dimensions in certain situations (e.g., vertical position vs. horizontal position). The relative importance of maintaining compatibility on each dimension is a function of how salient the dimensions are made by the task environment (Rubichi et al. 2006).

Stimulus–response compatibility effects occur for older adults as well as younger adults, with older adults typically showing larger compatibility effects that cannot be attributed entirely to general slowing (Proctor, Vu, and Pick 2005). Although older adults show a greater cost of incompatibility

than do younger adults, evidence indicates that the processing of information proceeds in a similar, though slower, manner (Vu and Proctor 2008). Because the older adults' response times increase disproportionately as a function of uncertainty, they benefit more from a precue that either indicates which of two tasks will be performed or reduces the number of possible stimulus and response alternatives (Vu and Proctor 2008). Implications of these findings for HCI are that older adults' performance will suffer more from incompatibility in designs, but this cost can be minimized by design strategies that limit the amount of information that must be processed.

Responses often produce effects in the environment, as, for example, when flipping a switch turns on a light. Studies have shown that speed of response selection is also influenced by such response-effect compatibility. Kunde (2001) had participants respond to the color of a single stimulus centered on a display screen by pressing one of four response keys, arranged in a row, with the index and middle fingers of the hands. Pressing a key filled in one box in a row of four outline boxes located above the response keys. Performance was faster when the mapping of keys to the filled-in boxes was spatially compatible than when it was incompatible. That response time is influenced by compatibility of a response with the effect that it produces implies that actions are selected and performed in anticipation of their consequences.

Because situations in which compatibility effects will influence performance are not always obvious, interface designers may make poor decisions if they rely only on their intuitions. Payne (1995), Vu and Proctor (2003), and Tlauka (2004) showed that naïve subjects can predict basic compatibility effects such as that performance will be better with a mapping that is spatially compatible than with one that is not. However, they do not accurately predict many other compatibility effects that occur such as the benefit of maintaining a consistent stimulus-response mapping rule. One encouraging finding is that estimates of relative compatibility can be improved by a small amount of experience performing with the different stimulus-response mappings (Vu and Proctor 2003). Designers need to be aware of the potential problems created by various types of incompatibility between display and response elements because their influences are not always obvious. A designer can get a better feel for the relative compatibility of alternative arrangements by performing tasks that use them. However, after the designer selects a few arrangements that would seem to yield good performance, more thorough usability testing of the remaining arrangements on groups of users needs to be performed.

2.4.3 RESPONSE EXECUTION

Motor programming refers to specification of the physical response that is to be made. This process is affected by variables such as relative stimulus-response frequency and movement direction. One factor that influences this stage is movement complexity. The longer the sequence of movements that is to be made upon occurrence of a stimulus in a choice-reaction task, the longer the RT to initiate the sequence

(Sternberg et al. 1978). This effect is thought to be due to the time required to load the movement sequence into a buffer before initiating the movements. Time to initiate the movement sequence decreases with practice, and fMRI evidence suggests that this decrease in RT involves distinct neural systems that support visuomotor learning of finger sequences and spatial learning of the locations of the finger movements on a keypad (Parsons, Harrington, and Rao 2005).

One of the most widely known relations attributed to response execution is Fitts's law, which describes the time to make aimed movements to a target location (Fitts 1954). This law, as originally specified by Fitts, is as follows:

$$\text{Movement Time} = a + b \log_2 (2D / W) \quad (2.2)$$

where a and b are constants, D is distance to the target, and W is target width. However, there are slightly different versions of the law. According to Fitts's law, movement time is a direct function of distance and an inverse function of target width. Fitts's law has been found to provide an accurate description of movement time in many situations, although alternatives have been proposed for certain situations. One factor that contributes to the increase in movement time as the index of difficulty increases is the need to make a corrective submovement based on feedback in order to hit the target location (Meyer et al. 1988).

The importance of Fitts's law for HCI is illustrated by the fact that the December 2004 issue of the *International Journal of Human-Computer Studies* was devoted to the fiftieth anniversary of Fitts's original study. In the preface to the issue, the editors, Guiard and Beudouin-Lafon (2004), state, "What has come to be known as Fitts's law has proven highly applicable in Human-Computer Interaction (HCI), making it possible to predict reliably the minimum time for a person in a pointing task to reach a specified target" (p. 747). Several illustrations of this point follow.

One implication of the law for interface design is that the slope of the function, b , may vary across different control devices, in which case, movement times will be faster for the devices that yield lower slopes. Card, English, and Burr (1978) conducted a study that evaluated how efficient text keys, step keys, a mouse, and a joystick are at a text-selection task, in which users selected text by positioning the cursor on the desired area and pressing a button or key. They showed that the mouse was the most efficient device for this task: Positioning time for the mouse and joystick could be accounted for by Fitts's law, with the slope of the function being less steep for the mouse; positioning time with the keys was proportional to the number of key strokes that had to be executed.

Another implication of Fitts's law is that any reduction in the index of difficulty should decrease the time for movements. Walker, Smelcer, and Nilsen (1991) evaluated movement time and accuracy of menu selection for the mouse. Their results showed that reducing the distance to be traveled (which reduces the index of difficulty) by placing the initial cursor in the middle of the menu, rather than the top, improved movement time. Placing a border around

the menu item in which a click would still activate that item, and increasing the width of the border as the travel distance increases, also improved performance. The reduction in movement time by use of borders is predicted by Fitts's law because borders increase the size of the target area. McGuffin and Balakrishnan (2005) showed that a similar reduction in movement time can be accomplished by expanding the target size while the movement is taking place.

Gillan et al. (1992) noted that designers must be cautious when applying Fitts's law to HCI because factors other than distance and target size play a role when using a mouse. Specifically, they proposed that the critical factors in pointing and dragging are different than those in pointing and clicking (which was the main task in Card, English, and Burr [1978] study). Gillan et al. showed that, for a text-selection task, both point-click and point-drag movement times can be accounted for by Fitts's law. For point-click sequences, the diagonal distance across the text object, rather than the horizontal distance, provided the best fit for pointing time. For point-drag, the vertical distance of the text provided the best fit. The reason why the horizontal distance is irrelevant is that the cursor must be positioned at the beginning of the string for the point-drag sequence. Thus, task requirements should be considered before applying Fitts's law to the interface design.

Motor adjustment deals with the transition from a central motor program to peripheral motor activity. Studies of motor adjustment have focused on the influence of foreperiod duration on motor preparation. In a typical study, a neutral warning signal is presented at various intervals before the onset of the imperative stimulus. Bertelson (1967) varied the duration of the warning foreperiod and found that RT reached a minimum at a foreperiod of 150 ms and then increased slightly at 200- and 300-ms foreperiods. However, error rate increased to a maximum at the 150-ms foreperiod and decreased slightly at the longer foreperiods. This relatively typical pattern suggests that it takes time to attain a state of high motor preparation, and that this state reflects an increased readiness to respond quickly at the expense of accuracy.

2.5 MEMORY IN INFORMATION PROCESSING

Memory refers to explicit recollection of information in the absence of the original stimulus and to persisting effects of that information on information processing that may be implicit. Memory may involve recall of an immediately preceding event or one many years in the past, knowledge derived from everyday life experiences and education, or procedures learned to accomplish complex perceptual-motor tasks. Memory can be classified into several categories. Episodic memory refers to memory for a specific event such as going to the movie last night, whereas semantic memory refers to general knowledge such as what a movie is. Declarative memory is verbalizable knowledge, and procedural memory is knowledge that can be expressed nonverbally. In other words, declarative memory is knowing that something is the case, whereas procedural memory is knowing how to

do something. For example, telling your friend your new phone number involves declarative memory, whereas riding a bicycle involves procedural knowledge. A memory test is regarded as explicit if a person is asked to judge whether a specific item or event has occurred before in a particular context; the test is implicit if the person is to make a judgment, such as whether a string of letters is a word or nonword, that can be made without reference to earlier "priming" events. In this section, we focus primarily on explicit episodic memory.

Three types of memory systems are customarily distinguished: sensory stores, short-term memory (STM; or working memory), and long-term memory (LTM). Sensory stores, which we will not discuss in detail, refer to brief modality-specific persistence of a sensory stimulus from which information can be retrieved for 1 or 2 seconds (see Nairne 2003). STM and LTM are the main categories by which investigations of episodic memory are classified, and as the terms imply, the distinction is based primarily on duration. The dominant view is that these are distinct systems that operate according to different principles, but there has been debate over whether the processes involved in these two types of memories are the same or different. An fMRI study by Talmi et al. (2005) found that recognition of early items in the list was accompanied by activation of areas in the brain associated with LTM, whereas recognition of recent items did not, supporting a distinction between STM and LTM stores.

2.5.1 SHORT-TERM (WORKING) MEMORY

STM refers to representations that are currently being used or have recently been used and last for a short duration. A distinguishing characteristic is that STM is of limited capacity. This point was emphasized in Miller's (1956) classic article, "The Magical Number Seven Plus or Minus Two," in which he indicated that capacity is not simply a function of the number of items, but rather the number of "chunks." For example, "i, b, m" are three letters, but most people can combine them to form one meaningful chunk of "IBM." Subsequent evidence indicates that the capacity of STM for verbal material is less than originally estimated by Miller, being three chunks when covert rehearsal is prevented (Chen and Cowan 2009). As a consequence of chunking, memory span is similar for strings of unrelated letters and strings of meaningful acronyms or words. Researchers refer to the number of items that can be recalled correctly, in order, as memory span. When rehearsal is not prevented, the memory span for words varies as a function of word length: The number of words that can be retained decreases as word length increases (Baddeley, Thomson, and Buchanan 1975). Evidence has indicated that the capacity is the number of syllables that can be said in about 2 seconds (Schweickert and Boruf 1986).

As most people are aware from personal experience, if distracted by another activity, information in STM can be forgotten quickly. With respect to HCI, Oulasvirta and Saariluoma (2004) note that diversion of attention from the current task to a competing task is a common occurrence, for example, when an unrequested pop-up dialog box requiring

an action appears on the screen. Laboratory experiments have shown that recall of a string of letters that is within the memory span decreases to close to chance levels over a retention interval of 18 seconds when rehearsal is prevented by an unrelated distractor task (Brown 1958; Peterson and Peterson 1959). This short-term forgetting was thought initially to be a consequence of decay of the memory trace due to prevention of rehearsal. However, Keppel and Underwood (1962) showed that proactive interference from items on previous lists is a significant contributor to forgetting. They found no forgetting at long retention intervals when only the first list in a series was examined, with the amount of forgetting being much larger for the second and third lists as proactive interference built up. Consistent with this interpretation, “release” from proactive inhibition, that is, improved recall, occurs when the category of the to-be-remembered items on the current list differs from that of previous lists (Wickens 1970).

As the complexity of an HCI task increases, one consequence is to overload STM. Jacko and Ward (1996) varied four different determinants of task complexity (multiple paths, multiple outcomes, conflicting interdependence among paths, or uncertain or probabilistic linkages) in a task requiring use of a hierarchical menu to acquire specified information. When one determinant was present, performance was slowed by approximately 50%, and when two determinants were present in combination, performance was slowed further. That is, as the number of complexity determinants in the interface increased, performance decreased. Jacko and Ward attributed the decrease in performance for all four determinants to the increased STM load they imposed.

The best-known model of STM is Baddeley and Hitch’s (1974) working memory model, which partitions STM into three main parts: central executive, phonological loop, and visuospatial sketchpad. The central executive is closely tied to the focus of attention. It is involved in computational processing, as in performing mental arithmetic, as well as in controlling and coordinating the actions of the phonological loop and visuospatial sketchpad. The phonological loop is composed of a phonological store that is responsible for storage of the to-be-remembered items, and an articulatory control process that is responsible for recoding verbal items into a phonological form and rehearsal of those items. The items stored in the phonological store decay over a short interval and can be refreshed through rehearsal from the articulatory control process. The visuospatial sketchpad retains information regarding visual and spatial information, and it is involved in mental imagery.

The working memory model has been successful in explaining several phenomena of STM (Baddeley 2000; 2003). However, the model cannot explain why memory span for visually presented material is only slightly reduced when subjects engage in concurrent articulatory suppression (such as saying the words “the” aloud repeatedly). Articulatory suppression should monopolize the phonological loop, preventing any visual items from entering it. To account for such findings, Baddeley revised the working memory model to include an episodic buffer (see Figure 2.3). The buffer is a




FIGURE 2.3 Baddeley’s (2000) revised working memory model. (Reprinted from *Trends in Cogn Sci*, 4, Baddeley, A. D., The episodic buffer: A new component of working memory? 421, Copyright 2000, with permission from Elsevier.)

limited capacity temporary store that can integrate information from the phonological loop, visuospatial sketchpad, and LTM. By attending to a given source of information in the episodic buffer, the central executive can create new cognitive representations that might be useful in problem solving.

2.5.2 LONG-TERM MEMORY

LTM refers to representations that can be remembered for durations longer than can be attributed to STM. LTM can involve information presented minutes ago or years ago. Initially, it was thought that the probability of an item being encoded into LTM was a direct function of the amount of time that it was in STM, or how much it was rehearsed. However, Craik and Watkins (1973) showed that rehearsal in itself is not sufficient, but rather that deep-level processing of the meaning of the material is the important factor in transferring items to LTM. They presented subjects with a list of words and instructed them that when the experimenter stopped the presentation, they were to recall the last word starting with the letter “a.” The number of other words between instances of “a” words was varied with the idea that the amount of time a word was rehearsed would depend on the number of words before the next “a” word. At the end of the session, subjects were given a surprise test in which they were to recall all “a” words. There was no effect of number of intervening words on recall, suggesting that although subjects rehearsed the words longer, their recall did not improve because the words were not processed deeply.

Craik and Watkins’ (1973) results are consistent with the levels of processing framework proposed by Craik and Lockhart (1972). According to this view, encoding proceeds in a series of analyses, from shallow perceptual features to deeper, semantic levels. The deeper the level of processing, the more strongly the item is encoded in memory. A key study supporting the levels of processing view is that of Hyde and Jenkins (1973). In their study, groups of subjects were presented a list of words for which they engaged in shallow processing (e.g., deciding whether each word contained a capital letter) or deep processing of it (e.g., identifying whether each word was a verb or a noun). Subjects were not told in advance that they would be asked to recall the words, but were given a surprise recall test at

the end of the session. Results showed that the deep processing group recalled more words than the shallow processing group. Of direct relevance to HCI, Oulasvirta, Kärkkäinen, and Laarni (2005) found that participants who viewed the content area of a web page had no better memory for the material than that guessed by a control group who had never seen the page, because the participants' task was to locate links on the page and not to process the content information.

Another well-known principle for LTM is encoding specificity, which states that the probability that a retrieval cue results in recollection of an earlier event is an increasing function of the match between the features encoded initially and those provided by the retrieval cue (Surprenant and Neath 2009). An implication of this principle is that memory will be context dependent. Godden and Baddeley (1975) demonstrated a context-dependent memory effect by having divers learn a list of words on land or under water, and recall the words on land or under water. Recall was higher for the group who learned on land when the test took place on land than under water, and vice versa for the group who learned under water. A related principle is that of transfer appropriate processing (Morris, Bransford, and Franks 1977). Morris et al. showed that deep-level semantic judgments during study produced better performance than shallow rhyme judgments on a standard recognition memory test. However, when the memory test required decisions about whether the test words rhymed with studied words, the rhyme judgments led to better performance than the semantic judgments. Brain imaging evidence consistent with transfer appropriate processing was obtained by Park and Rugg (2008), who found that word and picture stimuli on a recognition memory test produced greater activity in brain regions associated with those stimulus modes when the original study stimuli were also presented in the same mode. Healy, Wohldman, and Bourne (2005) have proposed that encoding specificity and transfer appropriate processing can be incorporated within the single principle of procedural reinstatement: Retention will be evident to the extent that the procedures engaged in during study or training are reinstated at the retention test.

Research has confirmed that the levels of processing framework must accommodate the effects of the retention context, as captured by the above principles, to explain the effects of processing performed during encoding. Although levels-of-processing has a strong effect on accuracy of explicit recall and recognition, Jacoby and Dallas (1981) found no effect on an implicit memory test. Later studies have shown a robust effect of levels-of-processing on implicit tests similar to that obtained for recall and recognition if the test is based on conceptual cues, rather than perceptual cues (Lee 2008).

2.5.3 OTHER FACTORS AFFECTING RETRIEVAL OF EARLIER EVENTS

Memory researchers have studied many factors that influence long-term retention. Not surprisingly, episodic memory improves with repetition of items or events. Also, massed repetition (repeating the same item in a row) is less effective

than spaced repetition (repeating the same item with one or more intervening items). This benefit for spaced repetition, called the spacing effect or lag effect, is often attributed to two main factors. First, study time for the same items appearing in succession is less than study time for the same items appearing further apart. Second, when the items are studied over a longer period of time, there is an opportunity for the items to be associated with different cues that can aid recall later. The spacing or lag effect is widespread and occurs for both recall and recognition (Hintzman 1974). Bahrack and Hall (2005) noted that a similar spacing benefit is found for learning lists of items when practice sessions, each with test and learning phases, are separated by several days. They presented evidence that a large part of the spacing benefit in this case arises from individuals determining which study strategies are more effective at promoting long-term retention and then using those strategies more.

Another widely studied phenomenon is the generation effect, in which recall is better when subjects have to generate the to-be-remembered words rather than just studying the words as they are presented (Slamecka and Graf 1978). In a generation effect experiment, subjects are divided into two groups: read and generate. Each group receives a series of words, with each word spelled out completely for the read group and missing letters for the generate group. An example is as follows:

Read group: CAT; ELEPHANT; GRAPE; CAKE

Generate group: C _ T; E _ E _ H _ NT; G _ APE; CAK _

The typical results show that subjects in the generate group can recall more words than those in the read group. One application of the generation effect to HCI is that when a computer user needs a password for an account, the system should allow the user to generate the password rather than providing him or her with one because the user would be more likely to recall the generated password. The common method of proactive password generation, in which users are asked to generate a password that meets certain restrictions (e.g., contain an uppercase letter, a lowercase letter, a digit, etc.), is intended to result in more memorable and secure passwords (see, e.g., Vu et al. 2007).

Events that precede or follow an event of interest can interfere with recall of that event. The former is referred to as proactive interference, and was discussed in the section on STM, and the latter is referred to as retroactive interference. One area of research in which retroactive interference is of central concern is that of eyewitness testimony. Loftus and Palmer (1974) showed that subsequent events could distort a person's memory of an event that the person witnessed. Subjects were shown a sequence of events depicting a car accident. Subsequently, they were asked the question, "How fast were the cars going when they _____ each other." When the verb "contacted" was used, subjects estimated the speed to be 32 mph, and only one-tenth of them reported seeing broken glass. However, when the verb "smashed" was used, the estimated speed increased to 41 mph, and almost one-third of the subjects reported seeing

broken glass. Demonstrations like these indicate not only that retroactive interference can cause forgetting of events, but that it also can cause the memory of events to be changed. More recent research has shown that completely false memories can be implanted (see Roediger and McDermott 1995).

Mnemonic techniques can also be used to improve recall. The basic idea behind mnemonics is to connect the to-be-remembered material with an established organizational structure that can be easily accessible later on. Two widely used mnemonic techniques are the pegword method (Wood and Pratt 1987) and the method of loci (Verhaeghen and Marcoen 1996). In the pegword method, a familiar rhyme provides the organizational structure. A visual image is formed between each pegword in the rhyme and the associated target item. At recall, the rhyme is generated, and the associated items come to mind. For the method of loci, locations from a well-known place, such as your house, are associated with the to-be-remembered items. Although specific mnemonic techniques are limited in their usefulness, the basic ideas behind them (utilizing imagery, forming meaningful associations, and using consistent encoding and retrieval strategies) are of broad value for improving memory performance.

Vu et al. (2007) examined the effectiveness of a “first-letter” mnemonic technique to help users relate individual characters of a password to a structured sentence in order to aid recall at a later time. In one condition, Vu et al. had users generate a sentence and take the first letter of each word in the sentence to form a password; in another condition, users generated a sentence that also included a number and special character embedded into the sentence and resulting password. Passwords generated using the first-letter technique were more memorable when users did not have to embed a digit and special character into the sentence, but were more secure (i.e., more resistant to cracking) when the sentence and resulting password included the digit and special character. Thus, when it comes to memory and security of computer passwords, there seems to be a trade-off between memorability and security.

Two additional factors have shown recently to benefit retrieval from memory. The first is repeated testing of items (Karpicke and Roediger 2008). The retrieval practice engendered by testing seems to be far more beneficial than additional studying of the items. The second is the relation of the to-be-remembered items to adaptive function (Nairne and Pandeirad 2010). Several studies have found evidence that survival-related words are retained better than ones that are not related to that adaptive function. From results like these, Nairne and Pandeirada have concluded, “to maximize retention in basic and applied settings it is useful to develop encoding techniques that are congruent with the natural design of memory systems” (p. 381).

2.6 ATTENTION IN INFORMATION PROCESSING

Attention is increased awareness directed at a particular event or action to select it for increased processing. This processing may result in enhanced understanding of the event, improved

performance of an action, or better memory for the event. Attention allows us to filter out unnecessary information so that we can focus on a particular aspect that is relevant to our goals. Several significant information-processing models of attention have been proposed.

2.6.1 MODELS OF ATTENTION

In an influential study, Cherry (1953) presented different messages to each ear through headphones. Subjects were to repeat aloud one of the two messages while ignoring the other. When subsequently asked questions about the two messages, subjects were able to accurately describe the message to which they were attending but could not describe anything except physical characteristics, such as gender of the speaker, about the unattended message.

To account for such findings, Broadbent (1958) developed the filter theory, which assumes that the nervous system acts as a single-channel processor. According to filter theory, information is received in a preattentive temporary store and then is selectively filtered, based on physical features such as spatial location, to allow only one input to access the channel. Broadbent’s filter theory implies that the meaning of unattended messages is not identified, but later studies showed that the unattended message could be processed beyond the physical level, in at least some cases (Treisman 1964).

To accommodate the finding that meaning of an unattended message can influence performance, Treisman (1964) reformulated filter theory into what is called the filter-attenuation theory. According to attenuation theory, early selection by filtering still precedes stimulus identification, but the filter only attenuates the information on unattended channels. This attenuated signal may be sufficient to allow identification if the stimulus is one with a low-identification threshold, such as a person’s name or an expected event. Deutsch and Deutsch (1963) proposed that unattended stimuli are always identified and the bottleneck occurs in later processing, a view called late-selection theory. The difference between attenuation theory and late-selection theory is that the latter assumes that meaning is fully analyzed, whereas the former does not.

Lavie et al. (2004) have proposed a load theory of attention, which they claim “resolves the long-standing early versus late selection debate” (p. 339). Specifically, the load theory includes two selective attention mechanisms, a perceptual selection mechanism and a cognitive control mechanism. When perceptual load is high (i.e., great demands are placed on the perceptual system), the perceptual mechanism excludes irrelevant stimuli from being processed. When memory load is high, it is not possible to suppress irrelevant information at a cognitive level. In support of load theory, Lavie et al. showed that interference from distracting stimuli is reduced under conditions of high perceptual load but increased under conditions of high working memory load.

In divided attention tasks, a person must attend to multiple sources of information simultaneously. Kahneman (1973) proposed a unitary resource model that views attention as a

single resource that can be divided up among different tasks in different amounts, based on task demands and voluntary allocation strategies. Unitary resource models provided the impetus for dual-task methodologies, such as performance operating characteristics, and mental workload analyses that are used widely in HCI (Eberts 1994). The expectation is that multiple tasks should produce interference when their resource demands exceed the supply that is available.

Many studies have shown that it is easier to perform two tasks together when they use different stimulus or response modalities than when they use the same modalities. Performance is also better when one task is verbal and the other visuospatial than when they are the same type. These result patterns provide the basis for multiple resource models of attention such as that of Wickens (1984). According to multiple resource models, different attentional resources exist for different sensory-motor modalities and coding domains. Multiple resource theory captures the fact that multiple-task performance typically is better when the tasks use different input–output modes than when they use the same modes. However, it is often criticized as being too flexible because new resources can be proposed arbitrarily to fit any finding of specificity of interference (Navon 1984).

A widely used metaphor for visual attention is that of a spotlight that is presumed to direct attention to everything in its field (Posner and Cohen 1984). Direction of attention is not necessarily the same as the direction of gaze because the attentional spotlight can be directed independently of fixation. Studies show that when a location is cued as likely to contain a target stimulus, but then a probe stimulus is presented at another location, a spatial gradient surrounds the attended location such that items nearer to the focus of attention are processed more efficiently than those farther away from it (Yantis 2000). The movement of the attentional spotlight to a location can be triggered by two types of cues: exogenous and endogenous. An exogenous cue is an external event such as the abrupt onset of a stimulus at a peripheral location that involuntarily draws the attentional spotlight to its location. Exogenous cues produce rapid performance benefits, which dissipate quickly, for stimuli presented at the cued location. This is followed by a period in which performance is worse for stimuli at the cued location than for ones presented at the uncued location, a phenomenon called inhibition of return (Posner and Cohen 1984). An endogenous cue is typically a symbol such as a central arrowhead that must be identified before a voluntary shift in attention to the designated location can be made. The performance benefits for endogenous cues take longer to develop and are sustained for a longer period of time when the cues are relevant, indicating that their benefits are due to conscious control of the attentional spotlight (Klein and Shore 2000).

Attentional focus is needed to detect change, and once attention is allocated to the processing of an event, there is a period in which it cannot be allocated to the processing of another event. Change blindness is the inability to detect sometimes large changes in a visual display or scene (Simons and Ambinder 2005). It has been demonstrated in the flicker

task, in which one scene alternates with another and the presence versus absence of a distinctive feature such as an aircraft engine or building is not detected. Change blindness also occurs in natural settings when attention is diverted momentarily. A closely related phenomenon is that of the attentional blink (Martens and Wyble 2010). In this paradigm, there is rapid presentation of a sequence of displays of visual stimuli. When a target stimulus is detected in one display, the probability of detecting a second target stimulus presented within the next several displays is reduced dramatically. Martens and Wyble attribute the attentional blink to a deficit in consolidation of the second target into a conceptual working-memory representation due to processing capacity being devoted to consolidation of the first target and being unavailable for processing of the second. They note that this limitation may be linked to a mechanism of attentional control.

In a visual search task, subjects are to detect whether a target is present among distractors. Treisman and Gelade (1980) developed feature integration theory to explain the results from visual search studies. When the target is distinguished from the distractors by a basic feature such as color (feature search), RT and error rate often show little increase as the number of distractors increases. However, when two or more features must be combined to distinguish the target from distractors (conjunctive search), RT and error rate typically increase sharply as the number of distractors increases. To account for these results, feature integration theory assumes that basic features of stimuli are encoded into feature maps in parallel across the visual field at a preattentive stage. Feature search can be based on this preattentive stage because a “target-present” response requires only detection of the feature. The second stage involves focusing attention on a specific location and combining features that occupy the location into objects. Attention is required for conjunctive search because responses cannot be based on detection of a single feature. According to feature integration theory, performance in conjunctive search tasks decreases as the number of distractors increases because attention must be moved sequentially across the search field until a target is detected or all items present have been searched. Feature integration theory served to generate a large amount of research on visual search that showed, as typically the case, that the situation is not as simple as depicted by the theory. This has resulted in modifications of the theory, as well as alternative theories. For example, Wolfe’s (2007) Guided Search Theory maintains the distinction between an initial stage of feature maps and a second stage of attentional binding, but assumes that the second stage is guided by the initial feature analysis.

In HCI, a common visual search task involves locating menu items. When users know exactly what option to search for, identity matching can be used, in which users search the display for the menu name that they want to find. Perlman (1984) suggested that when identity search is used, the menu options should be displayed in alphabetical order to facilitate search. When users do not know where an option is included within a main list of menus, inclusion matching is used.

The users must decide within which group the specific option would be categorized and then search the list of items for that group. With inclusion matching, search times may be longer for items that can be classified in more than one of the main groupings or when the items are less well-known examples of a main grouping (Somberg and Picardi 1983). Equivalence search occurs when the users know what option to select, but does not know how that option is labeled. McDonald, Stone, and Liebelt (1983) showed that alphabetical and categorical organizations yield shorter search times than randomized organization for equivalence search. Search can also be affected by the breadth versus depth of the menu design. Lee and MacGregor (1985) showed that deep hierarchies are preferred over broad ones. However, more recently, Tullis, Tranquada, and Siegel (2011) suggested that, for complex or ambiguous situations, there is a benefit for broad menu designs because they facilitate comparison between categories. The main point is that when structuring menus, designers must consider the type of search in which the user would most likely be engaged.

The role of attention in response selection has been investigated extensively using the psychological refractory period (PRP) paradigm (Pashler 1998). In the PRP paradigm, a pair of choice-reaction tasks must be performed, and the stimulus onset asynchrony (SOA) of the second stimulus is presented at different intervals. RT for Task 2 is slowed at short SOAs, and this phenomenon is called the PRP effect. The experimental results have been interpreted with what is called locus of slack logic (Schweickert 1978), which is an extension of additive factors logic to dual-task performance. The basic idea is that if a Task 2 variable has its effect prior to a bottleneck, that variable will have an underadditive interaction with SOA. This underadditivity occurs because, at short SOAs, the slack period during which postbottleneck processing cannot begin can be used for continued processing for the more difficult condition. If a Task 2 variable has its effect after the bottleneck, the effect will be additive with SOA.

The most widely accepted account of the PRP effect is the response-selection bottleneck model (Pashler 1998). The primary evidence for this model is that perceptual variables typically have underadditive interactions with SOA, implying that their effects are before the bottleneck. In contrast, postperceptual variables typically have additive effects with SOA, implying that their effects are after the bottleneck. There has been dispute as to whether there is also a bottleneck at the later stage of response initiation (De Jong 1993), whether the response-selection bottleneck is better characterized as a parallel processor of limited capacity that divides resources among to-be-performed tasks (Tombu and Jolicoeur 2005), and whether the apparent response-selection bottleneck is structural or simply a strategy adopted by subjects to comply with task instructions (Meyer and Kieras 1997). This latter approach is consistent with an emphasis on the executive functions of attention in the coordination and control of cognitive processes (Monsell and Driver 2000).

2.6.2 AUTOMATICITY AND PRACTICE

Attention demands are high when a person first performs a new task. However, these demands decrease and performance improves as the task is practiced. Because the quality of performance and attentional requirements change substantially as a function of practice, it is customary to describe performance as progressing from an initial cognitively demanding phase to a phase in which processing is automatic (Anderson 1982; Fitts and Posner 1967).

With the largest benefits occurring early in practice, the time to perform virtually any task from choice RT to solving geometry problems decreases with practice. Newell and Rosenbloom (1981) proposed a power function to describe the changes in RT with practice:

$$RT = BN^{-\alpha} \quad (2.3)$$

where N is the number of practice trials, B is RT on the first trial, and α is the learning rate. Although the power function has become widely accepted as a law that describes the changes in RT, Heathcote, Brown, and Mewhort (2000) indicated that it does not fit the functions for individual performers adequately. They showed that exponential functions provided better fits than power functions to 40 individual data sets, and proposed a new exponential law of practice. The defining characteristic of the exponential function is that the relative learning rate is a constant at all levels of practice, whereas, for the power function, the relative learning rate is a hyperbolically decreasing function of practice trials.

2.7 PROBLEM SOLVING AND DECISION MAKING

Beginning with the work of Newell and Simon (1972), it has been customary to analyze problem solving in terms of a problem space. The problem space consists of the following: (1) an initial state, (2) a goal state that is to be achieved, (3) operators for transforming the problem from the initial state to the goal state in a sequence of steps, and (4) constraints on application of the operators that must be satisfied. The problem-solving process itself is conceived of as a search for a path that connects the initial and goal states.

Because the size of a problem space increases exponentially with the complexity of the problem, most problem spaces are well beyond the capacity of STM. Consequently, for problem solving to be effective, search must be constrained to a limited number of possible solutions. A common way to constrain search is through the use of heuristics. For example, people often use a means-ends heuristic for which at each step, an operator is chosen that will move the current state closer to the goal state (Atwood and Polson 1976). Such heuristics are called weak methods because they do not require much knowledge about the exact problem domain. Strong methods, such as those used by experts, rely on prior domain-specific knowledge and do not require much search because they are based on established principles applicable only to certain tasks.

The problem space must be an appropriate representation of the problem, if the problem is to be solved. One important method for obtaining an appropriate problem space is to use analogy or metaphor. Analogy enables a shift from a problem space that is inadequate to one that may allow the goal state to be reached. There are several steps in using analogies (Holland et al. 1986), including detecting similarity between source and target problems, and mapping the corresponding elements of the problems. Humans are good at mapping the problems, but poor at detecting that one problem is an analog of another. An implication for HCI is that potential analogs should be provided to users for situations in which they are confronted by novel problems.

The concept of mental model, which is closely related to that of the problem space, has become widely used in recent years (see Payne, this volume). The general idea of mental models with respect to HCI is that as the user interacts with the computer, she or he receives feedback from the system that allows him/her to develop a representation of how the system is functioning for a given task. The mental model incorporates the goals of the user, the actions taken to complete the goals, and expectations of the system's output in response to the actions. A designer can increase the usability of an interface by using metaphors that allow transfer of an appropriate mental model (e.g., the desktop metaphor), designing the interface to be consistent with other interfaces with which the user is familiar (e.g., the standard web interface), and conveying the system's functions to the user in a clear and accurate manner. Feedback to the user is perhaps the most effective way to communicate information to the user and can be used to guide the user's mental model about the system.

Humans often have to make choices for situations in which the outcome depends on events that are outside of their control. According to expected utility theory, a normative theory of decision making under uncertainty, the decision maker should determine the expected utility of a choice by multiplying the subjective utility of each outcome by the outcome's probability and summing the resulting values (Hastie and Dawes 2010). The expected utility should be computed for each choice, and the optimal decision is the choice with the highest expected utility. It should be clear from this description that for all but the simplest of problems, a human decision maker cannot operate in this manner. To do so would require attending to multiple cues that exceed attentional capacity, accurate estimates of probabilities of various events, and maintenance of, and operation on, large amounts of information that exceeds STM capacity.

Research of Kahneman and Tversky (2000) and others has shown that what people do when the outcome associated with a choice is uncertain is to rely heavily on decision-making heuristics. These heuristics include representativeness, availability, and anchoring. The representativeness heuristic is that the probability of an instance being a member of a particular category is judged on the basis of how representative the instance is of the category. The major limitation of the

representativeness heuristic is that it ignores base rate probabilities for the respective categories. The availability heuristic involves determining the probability of an event based on the ease with which instances of the event can be retrieved. The limitation is that availability is affected not only by relative frequency but also by other factors. The anchoring heuristic involves making a judgment regarding probabilities of alternative states based on initial information, and then adjusting these probabilities from this initial "anchor" as additional information is received. The limitation of anchoring is that the initial judgment can produce a bias for the probabilities. Although heuristics are useful, they may not always lead to the most favorable decision. Consequently, designers need to make sure that the choice desired for the user in a particular situation is one that is consistent with the user's heuristic biases.

2.8 SUMMARY AND CONCLUSION

The methods, theories, and models in human information processing are currently well developed. The knowledge in this area, of which we are only able to describe at a surface level in this chapter, is relevant to a wide range of concerns in HCI, from visual display design to representation and communication of knowledge. For HCI to be effective, the interaction must be made compatible with the human information-processing capabilities. Cognitive architectures that incorporate many of the facts about human information processing have been developed that can be applied to HCI. The Model Human Processor of Card, Moran, and Newell (1983) is the most widely known, but other more recent architectures, including the adaptive control of thought model of Anderson and colleagues (Anderson, Matessa, and Lebiere 1997), the State, Operator, and Result (SOAR) Model of Newell and colleagues (Howes and Young 1997), and the EPIC Model of Kieras and Meyer (1997) have considerable utility for the field, as demonstrated in Chapter 5 of this volume.

The human information-processing approach emphasizes laboratory research in which fundamental cognitive processes and principles thought to be of broad generalizability are established. Although this approach has been highly successful in many respects, some researchers think that more emphasis should be placed on real-world behavior in natural environments. Alternative approaches to perception, cognition, and action with such emphasis include the following. The ecological approach associated with Gibson (1979) places emphasis on analyzing the perceptual information that is available in the optic array and the dynamics of this information as the individual interacts with the environment. The cybernetic view, that cognition emerges as a consequence of motor control over sensory feedback, stresses self-regulated control of perception and cognition (Smith and Henning 2005). The situated cognition approach focuses on the need to understand behavior in specific contexts in which, for example, a computer application will be used (Kiekel and Cooke 2011). A recently popular approach

is that of embodied cognition, according to which knowledge is acquired and processed through interactions of the body with the environment (e.g., Sherman, Gangi, and White 2010). One common feature of these alternative accounts is an emphasis on the relation among perception, cognition, and action. We agree that, in certain areas of information-processing research, action has been viewed as a final stage that does not influence the prior stages of perception and cognition. However, in other areas, such as that of human performance, action has been emphasized since the earliest applications of the information-processing approach (Fitts and Posner 1967; since 1975, one division of the *Journal of Experimental Psychology* has been subtitled *Human Perception and Performance*). From our perspective, information-processing analyses and models will continue to be useful tools for understanding and predicting human behavior both in general and in HCI in particular.

REFERENCES

- Anderson, J. R. 1982. Acquisition of cognitive skill. *Psychol Rev* 89:369–406.
- Anderson, J. R., D. Bothell, M. D. Byrne, S. Douglass, C. Lebiere, and Y. Qin. 2004. An integrated theory of the mind. *Psychol Rev* 111:1036–60.
- Anderson, J. R., J. M. Fincham, and S. Douglass. 1999. Practice and retention: A unifying analysis. *J Exp Psychol Learn Mem Cogn* 25:1120–36.
- Anderson, J. R., M. Matessa, and C. Lebiere. 1997. ACT-R: A theory of higher level cognition and its relation to visual attention. *Hum Comput Interact* 12:439–62.
- Atwood, M. E., and P. G. Polson. 1976. A process model for water jug problems. *Cogn Psychol* 8:191–216.
- Baddeley, A. D. 2000. The episodic buffer: A new component of working memory? *Trends Cogn Sci* 4:421.
- Baddeley, A. D. 2003. Working memory and language: An overview. *J Commun Disord* 36:189–208.
- Baddeley, A. D., and G. J. Hitch. 1974. Working memory. In *The Psychology of Learning and Motivation*, ed. G. H. Bower, vol. 8, 47–89. New York: Academic Press.
- Baddeley, A. D., N. Thomson, and M. Buchanan. 1975. Word length and the structure of short-term memory. *J Verbal Learn Behav* 14:575–89.
- Bahrick, H. P., and L. K. Hall. 2005. The importance of retrieval failures to long-term retention: A metacognitive explanation of the spacing effect. *J Mem Lang* 52:566–77.
- Bertelson, P. 1967. The time course of preparation. *Q J Exp Psychol* 19:272–79.
- Boldini, A., R. Russo, and S. E. Avons. 2004. One process is not enough! A speed-accuracy tradeoff study of recognition memory. *Psychon Bull Rev* 11:353–61.
- Bower, G. H., and E. R. Hilgard. 1981. *Theories of Learning*. 5th ed. Englewood Cliffs, NJ: Prentice Hall.
- Broadbent, D. E. 1958. *Perception and Communication*. Oxford, UK: Pergamon Press.
- Brown, J. (1958). Some tests of the decay theory of immediate memory. *Q J Exp Psychol* 10:12–21.
- Bussemeyer, J. R., and A. Diederich. 2010. *Cognitive Modeling*. Thousand Oaks, CA: Sage.
- Card, S. K., W. K. English, and B. J. Burr. 1978. Evaluation of the mouse, rate-controlled isometric joystick, step keys, and text keys for text selection on a CRT. *Ergonomics* 21:601–13.
- Card, S. K., T. P. Moran, and A. Newell. 1983. *The Psychology of Human-Computer Interaction*. Hillsdale, NJ: Erlbaum.
- Chen, Z., and N. Cowan. 2009. Core verbal working-memory capacity: The limit in words retained without covert articulation. *Q J Exp Psychol* 62:1420–9.
- Cherry, E. C. 1953. Some experiments on the recognition of speech, with one and with two ears. *J Acoust Soc Am* 25:975–9.
- Craik, F. I. M., and R. S. Lockhart. 1972. Levels of processing: A framework for memory research. *J Verbal Learn Verbal Behav* 11:671–84.
- Craik, F. I. M., and M. J. Watkins. 1973. The role of rehearsal in short-term memory. *J Verbal Learn Verbal Behav* 12:599–607.
- De Jong, R. 1993. Multiple bottlenecks in overlapping task performance. *J Exp Psychol Hum Percept Perform* 19:965–80.
- Deutsch, J. A., and D. Deutsch. 1963. Attention: Some theoretical considerations. *Psychol Rev* 70:80–90.
- Donders, F. C. 1868/1969. On the speed of mental processes. In *Acta Psychologica*, 30, *Attention and Performance II*, ed. W. G. Koster, 412–31. Amsterdam, Netherlands: North-Holland.
- Duncan, J. 1977. Response selection rules in spatial choice reaction tasks. In *Attention and performance VI*, ed. S. Dornic, 49–71. Hillsdale, NJ: Erlbaum.
- Eberts, R. E. 1994. *User Interface Design*. Englewood Cliffs, NJ: Prentice Hall.
- Eimer, M. 1998. The lateralized readiness potential as an on-line measure of central response activation processes. *Behav Res Methods Instrum Comput* 30:146–56.
- Fitts, P. M. 1951. Engineering psychology and equipment design. In *Handbook of Experimental Psychology*, ed. S. S. Stevens, 1287–340. New York: Wiley.
- Fitts, P. M. 1954. The information capacity of the human motor system in controlling the amplitude of movement. *J Exp Psychol* 47:381–91.
- Fitts, P. M., and R. L. Deininger. 1954. S-R compatibility: Correspondence among paired elements within stimulus and response codes. *J Exp Psychol* 48:483–92.
- Fitts, P. M., and M. I. Posner. 1967. *Human Performance*. Belmont, CA: Brooks/Cole.
- Fitts, P. M., and C. M. Seeger. 1953. S-R compatibility: Spatial characteristics of stimulus and response codes. *J Exp Psychol* 46:199–210.
- Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Boston, MA: Houghton Mifflin.
- Gillan, D. J., K. Holden, S. Adam, M. Rudisill, and L. Magee. 1992. How should Fitts's law be applied to human-computer interaction? *Interact Comput* 4:291–313.
- Godden, D. R., and A. D. Baddeley. 1975. Context-dependent memory in two natural environments: On land and underwater. *Br J Psychol* 66:325–31.
- Gonzalez, C., F. J. Lerch, and C. Lebiere. 2003. Instance-based learning in real-time dynamic decision making. *Cogn Sci* 27:591–635.
- Gobelny, J., W. Karwowski, and C. Drury. 2005. Usability of graphical icons in the design of human-computer interfaces. *Int J Hum Comput Interact* 18:167–82.
- Guiard, Y., and M. Beaudouin-Lafon. 2004. Fitts's law 50 years later: Applications and contributions from human-computer interaction. *Int J Hum Comput Stud* 61:747–50.
- Hastie, R., and R. M. Dawes. 2010. *Rational Choice in An Uncertain World: The Psychology of Judgment and Decision Making*. Thousand Oaks, CA: Sage.
- Healy, A. F., E. L. Wohlmann, and L. E. Bourne Jr. 2005. The procedural reinstatement principle: Studies on training, retention, and transfer. In *Experimental Cognitive Psychology and Its Applications*, ed. A. F. Healy, 59–71. Washington, DC: American Psychological Association.

- Heathcote, A., S. Brown, and D. J. K. Mewhort. 2000. The power law repealed: The case for an exponential law of practice. *Psychono Bull Rev* 7:185–207.
- Hick, W. E. 1952. On the rate of gain of information. *Q J Exp Psychol* 4:11–26.
- Hintzman, D. L. 1974. Theoretical implications of the spacing effect. In *Theories of Cognitive Psychology: The Loyola Symposium*, ed. R. L. Solso, 77–99. Hillsdale, NJ: Erlbaum.
- Hoffman, R. R., and K. A. Deffenbacher. 1992. A brief history of applied cognitive psychology. *Appl Cogn Psychol* 6:1–48.
- Holland, J. H., K. J. Holyoak, R. E. Nisbett, and P. R. Thagard. 1986. *Induction*. Cambridge, MA: MIT Press.
- Hommel, B., and W. Prinz, eds. 1997. *Theoretical Issues in Stimulus-Response Compatibility*. Amsterdam, Netherlands: North-Holland.
- Howes, A., and R. M. Young. 1997. The role of cognitive architecture in modeling the user: Soar's learning mechanism. *Hum Comput Interact* 12:311–43.
- Huettel, S. A., A. W. Song, and G. McCarthy. 2004. *Functional Magnetic Resonance Imaging*. Sunderland, MA: Sinauer Associates.
- Hyde, T. S., and J. J. Jenkins. 1973. Recall of words as a function of semantic, graphic, and syntactic orienting tasks. *J Verbal Learn Verbal Behav* 12:471–80.
- Hyman, R. 1953. Stimulus information as a determinant of reaction time. *J Exp Psychol* 45:188–96.
- Jacko, J. A., and K. G. Ward. 1996. Toward establishing a link between psychomotor task complexity and human information processing. In *19th International Conference on Computers and Industrial Engineering*, vol. 31, 533–6.
- Jacoby, L. L., and M. Dallas. 1981. On the relationship between autobiographical memory and perceptual learning. *J Exp Psychol Gen* 110:306–40.
- Kahneman, D. 1973. *Attention and Effort*. Englewood Cliffs, NJ: Prentice Hall.
- Kahneman, D., and A. Tversky, eds. 2000. *Choices, Values, and Frames*. New York: Cambridge University Press.
- Karpicke, J. D., and H. L. Roediger III. 2008. The critical importance of retrieval for learning. *Science* 319:966–8.
- Kauhanen, L., P. Jylänki, J. Lehtonen, P. Rantanen, H. Alaranta, and M. Sams. 2007. EEG-based brain-computer interface for tetraplegics. *Compu Intell Neurosci* Volume 2007, Article ID 23864, 11 pages doi:10.1155/2007/23864.
- Keppel, G., and B. J. Underwood. 1962. Proactive inhibition in short-term retention of single items. *J Verbal Learn Verbal Behav* 1:153–61.
- Kiekel, P. A., and N. J. Cooke. 2011. Human factors aspects of team cognition. In *Handbook of Human Factors in Web Design*, 2nd ed., ed. K.-P. L. Vu and R. W. Proctor, 107–23. Boca Raton, FL: CRC Press.
- Kieras, D. E., and D. E. Meyer. 1997. An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. *Hum Comput Interact* 12:391–438.
- Klein, R. M., and D. I. Shore. 2000. Relation among modes of visual orienting. In *Control of Cognitive Processes: Attention and Performance XVIII*, ed. S. Monsell and J. Driver, 195–208. Cambridge, MA: MIT Press.
- Kornblum, S., T. Hasbroucq, and A. Osman. 1990. Dimensional overlap: Cognitive basis for stimulus-response compatibility—A model and taxonomy. *Psychol Rev* 97:253–70.
- Kunde, W. 2001. Response-effect compatibility in manual choice reaction tasks. *J Exp Psychol Hum Percept Perform* 27:387–94.
- Lachman, R., J. L. Lachman, and E. C. Butterfield. 1979. *Cognitive Psychology and Information Processing: An Introduction*. Hillsdale, NJ: Erlbaum.
- Lavie, N., A. Hirst, J. W. de Fockert, and E. Viding. 2004. Load theory of selective attention and cognitive control. *J Exp Psychol Gen* 133:339–54.
- Lee, Y.-S. 2008. Levels-of-processing effects on conceptual automatic memory. *Eur J Cogn Psychol* 20:936–54.
- Lee, E., and J. MacGregor. 1985. Minimizing user search time in menu retrieval systems. *Hum Factors* 27:157–62.
- Lees, M. N., J. D. Cosman, J. D. Lee, M. Rizzo, and N. Fricke. 2010. Translating cognitive neuroscience to the driver's operational environment: A neuroergonomics approach. *Am J Psychol* 123:391–411.
- Loftus, E. F., and J. C. Palmer. 1974. Reconstruction of automobile destruction: An example of the interaction between language and memory. *J Exp Psychol Hum Learn Mem* 4:19–41.
- Lu, C.-H., and R. W. Proctor. 1995. The influence of irrelevant location information on performance: A review of the Simon effect and spatial Stroop effects. *Psychono Bull Rev* 2:174–207.
- MacKenzie, I. S., and R. W. Soukoreff. 2002. Text entry for mobile computing: Models and methods, theory and practice. *Hum Comput Interact* 17:147–98.
- Mackworth, N. H. 1948. The breakdown of vigilance during prolonged visual search. *Q J Exp Psychol* 1:6–21.
- Macmillan, N. A., and C. D. Creelman. 2005. *Detection Theory: A User's Guide*. 2nd ed. Mahwah, NJ: Erlbaum.
- Martens, S., and B. Wyble. 2010. The attentional blink: Past, present, and future of a blind spot in perceptual awareness. *Neurosci Biobehav Rev* 34:947–57.
- McCarley, J. S. 2009. Effects of speed-accuracy instructions on oculomotor scanning and target recognition in a simulated baggage X-ray screening task. *Ergonomics* 52:325–33.
- McClelland, J. L. 1979. On the time relations of mental processes: A framework for analyzing processes in cascade. *Psychol Rev* 88:375–407.
- McDonald, J. E., J. D. Stone, and L. S. Liebelt. 1983. Searching for items in menus: The effects of organization and type of target. In *Proceedings of the Human Factors Society 27th Annual Meeting*, 289–338. Hillsdale, NJ: Erlbaum.
- McGuffin, M. J., and R. Balakrishnan. 2005. Fitts's law and expanding targets: Experimental studies and designs for user interfaces. *ACM Trans Comput Hum Interact* 12:388–422.
- Melara, R. D., H. Wang, K.-P. L. Vu, and R. W. Proctor. 2008. Attentional origins of the Simon effect: Behavioral and electrophysiological evidence. *Brain Res* 1215:147–59.
- Meyer, D. E., R. A. Abrams, S. Kornblum, C. E. Wright, and J. E. K. Smith. 1988. Optimality in human motor performance: Ideal control of rapid aimed movements. *Psychol Rev* 86:340–70.
- Meyer, D. E., and D. E. Kieras. 1997. A computational theory of executive cognitive processes and multiple-task performance: Part 2. Accounts of psychological refractory-period phenomena. *Psychol Rev* 104:749–91.
- Miller, G. A. 1956. The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychol Rev* 63:81–97.
- Miller, J. 1988. Discrete and continuous models of human information processing: Theoretical distinctions and empirical results. *Acta Psychol* 67:191–257.
- Monsell, S., and J. Driver, eds. 2000. *Control of Cognitive Processes: Attention and Performance XVIII*. Cambridge, MA: MIT Press.

- Morris, C. D., J. D. Bransford, and J. J. Franks. 1977. Levels of processing versus transfer appropriate processing. *J Verbal Learn Verbal Behav* 16:519–33.
- Nairne, J. S. 2003. Sensory and working memory. In *Experimental Psychology*, vol. 4 of *Handbook of Psychology*, ed. A. F. Healy and R. W. Proctor. Editor-in-Chief: I. B. Weiner. Hoboken, NJ: Wiley.
- Nairne, J. S., and Pandeirada, J. N. S. 2010. Adaptive memory: Nature's criterion and the functionalist agenda. *Am J Psychol* 123:381–90.
- Navon, D. 1984. Resources—a theoretical soup stone? *Psychol Rev* 91:216–34.
- Newell, A., and P. S. Rosenbloom. 1981. Mechanisms of skill acquisition and the law of practice. In *Cognitive Skills and Their Acquisition*, ed. J. R. Anderson, 1–55. Hillsdale, NJ: Erlbaum.
- Newell, A., and H. A. Simon. 1972. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice Hall.
- Oulasvirta, A., L. Kärkkäinen, and J. Laarni. 2005. Expectations and memory in link search. *Comput Hum Behav* 21:773–89.
- Oulasvirta, A., and P. Saariluoma. 2004. Long-term working memory and interrupting messages in human-computer interaction. *Behav Inf Technol* 23:53–64.
- Pachella, R. G. 1974. The interpretation of reaction time in information-processing research. In *Human Information Processing: Tutorials in Performance and Cognition*, ed. B. H. Kantowitz, 41–82. Hillsdale, NJ: Erlbaum.
- Parasuraman, R. 2003. Neuroergonomics: Research and practice. *Theor Issues Ergon Sci* 4:5–20.
- Parasuraman, R., and D. R. Davies. 1977. A taxonomic analysis of vigilance performance. In *Vigilance: Theory, Operational Performance, and Physiological Correlates*, ed. R. R. Mackie, 559–74. New York: Plenum.
- Park, H., and M. D. Rugg. 2008. The relationship between study processing and the effects of cue congruency at retrieval: fMRI support for transfer appropriate processing. *Cereb Cortex* 18:868–75.
- Parsons, M. W., D. L. Harrington, and S. M. Rao. 2005. Distinct neural systems underlie learning visuomotor and spatial representations of motor skills. *Hum Brain Mapp* 24:229–47.
- Pashler, H. 1998. *The Psychology of Attention*. Cambridge, MA: MIT Press.
- Payne, S. J. 1995. Naïve judgments of stimulus-response compatibility. *Hum Factors* 37:495–506.
- Percival, L. C., and T. K. Noonan. 1987. Computer network operation: Applicability of the vigilance paradigm to key tasks. *Hum Factors* 29:685–94.
- Perlman, G. 1984. Making the right choices with menus. *Proceedings of INTERACT '84*, 291–5. London, UK: IFIP.
- Peterson, L. R., and M. J. Peterson. 1959. Short-term retention of individual verbal items. *J Exp Psychol* 58:193–8.
- Posner, M. I. 1986. Overview. In *Handbook of Perception and Human Performance*, vol. 2 of *Cognitive Processes and Performance*, ed. K. R. Boff, L. Kaufman, and J. P. Thomas, V3–10. New York: Wiley.
- Posner, M. I., and Y. Cohen. 1984. Components of visual orienting. In *Attention and Performance X*, ed. H. Bouma, and D. G. Bouwhuis, 531–56. Hillsdale, NJ: Erlbaum.
- Proctor, R. W., and Y. S. Cho. 2006. Polarity correspondence: A general principle for performance of speeded binary classification tasks. *Psychol Bull* 132:416–42.
- Proctor, R. W., and T. Van Zandt. 2008. *Human Factors in Simple and Complex Systems*. 2nd ed. Boca Raton, FL: CRC Press.
- Proctor, R. W., and K.-P. L. Vu. 2006. *Stimulus-Response Compatibility: Data, Theory and Application*. Boca Raton, FL: CRC Press.
- Proctor, R. W., and K.-P. L. Vu. 2010. Cumulative knowledge and progress in human factors. *Annu Rev Psychol* 61:623–51.
- Proctor, R. W., K.-P. L. Vu, and D. F. Pick. 2005. Aging and response selection in spatial choice tasks. *Hum Factors* 47:250–70.
- Ratcliff, R., and P. L. Smith. 2004. A comparison of sequential sampling models for two-choice reaction time. *Psychol Rev* 111:333–67.
- Reeve, T. G., and R. W. Proctor. 1990. The salient features coding principle for spatial- and symbolic-compatibility effects. In *Stimulus-Response Compatibility: An Integrated Perspective*, ed. R. W. Proctor and T. G. Reeve, 163–80. Amsterdam, Netherlands: North-Holland.
- Roediger III, H. L., and K. B. McDermott. 1995. Creating false memories: Remembering words not presented in lists. *J Exp Psychol Learn Mem Cogn* 21:803–14.
- Roscoe, S. 2011. Historical overview of human factors and ergonomics. In *Handbook of Human Factors in Web Design*, 2nd ed., ed. K.-P. L. Vu and R. W. Proctor, 3–12. Boca Raton: CRC Press.
- Rotello, C. M., and N. A. Macmillan. 2006. Remember-know models as decision strategies in two experimental paradigms. *J Mem Lang* 55:479–94.
- Rubichi, S., K.-P. L. Vu, R. Nicoletti, and R. W. Proctor. 2006. Spatial coding in two dimensions. *Psychon Bull Rev* 13:201–16.
- Rugg, M. D., and M. G. H. Coles, eds. 1995. *Electrophysiology of Mind: Event-Related Brain Potentials and Cognition*. Oxford, England UK: Oxford University Press.
- Salthouse, T. A. 2005. From description to explanation in cognitive aging. In *Cognition and Intelligence: Identifying the Mechanisms of the Mind*, ed. R. J. Sternberg and J. E. Pretz, 288–305. Cambridge: Cambridge University Press.
- Sanders, A. F. 1998. *Elements of Human Performance*. Mahwah, NJ: Erlbaum.
- Sartori, G., and C. Umiltà. 2000. How to avoid the fallacies of cognitive subtraction in brain imaging. *Brain Lang* 74:191–212.
- Schweickert, R. 1978. A critical path generalization of the additive factor method: Analysis of a Stroop task. *J Math Psychol* 18:105–39.
- Schweickert, R., and B. Boruff. 1986. Short-term memory capacity: Magic number or magic spell? *J Exp Psychol Learn Mem Cogn* 12:419–25.
- Shaffer, L. H. 1965. Choice reaction with variable S-R mapping. *J Exp Psychol* 70:284–8.
- Shepard, R. N., and J. Metzler. 1971. Mental rotation of three-dimensional objects. *Science* 171:701–3.
- Sherman, D. K., C. Gangi, and M. L. White. 2010. Embodied cognition and health persuasion: Facilitating intention-behavior consistency via motor manipulations. *J Exp Soc Psychol* 46:461–4.
- Simon, J. R. 1990. The effects of an irrelevant directional cue on human information processing. In *Stimulus-Response Compatibility: An Integrated Perspective*, ed. R. W. Proctor and T. G. Reeve, 31–86. Amsterdam: North-Holland.
- Simons, D. J., and M. S. Ambinder. 2005. Change blindness: Theory and consequences. *Curr Dir Psychol Sci* 14:44–8.
- Slamecka, N. J., and P. Graf. 1978. The generation effect: Delineation of a phenomenon. *J Exp Psychol Hum Learn Mem* 4:592–604.

- Smith, T. J., and R. A. Henning. 2005. Cybernetics of augmented cognition as an alternative to information processing. In *Proceedings of HCI International 2005*. Mahwah, NJ: Erlbaum.
- Somberg, B. L., and M. C. Picardi. 1983. Locus of information familiarity effect in search of computer menus. In *Proceedings of the Human Factors Society 27th Annual Meeting*, 826–30. Santa Monica, CA: Human Factors Society.
- Sternberg, S. 1969. The discovery of processing stages: Extensions of Donders' method. In *Attention and Performance II Acta Psychologica*, ed. W. G. Koster, vol. 30, 276–315. North Holland: Amsterdam.
- Sternberg, S., S. Monsell, R. L. Knoll, and C. E. Wright. 1978. The latency and duration of rapid movement sequences. In *Information Processing in Motor Control and Learning*, ed. G. E. Stelmach, 117–52. New York: Academic Press.
- Stevenson, R. A., S. Kim, and T. W. James. 2009. An additive-factors design to disambiguate neuronal and areal convergence: Measuring multisensory interactions between audio, visual, and haptic sensory streams using fMRI. *Exp Brain Res* 198:183–94.
- Surprenant, A. M., and I. Neath. 2009. *Principles of Memory*. Philadelphia, PA: Psychology Press.
- Talmi, D., C. L. Grady, Y. Goshen-Gottstein, and M. Moscovitch. 2005. Neuroimaging the serial position curve: A test of single-store versus dual-store models. *Psychol Sci* 16:716–23.
- Teichner, W. H., and M. J. Krebs. 1974. Laws of visual choice reaction time. *Psychol Rev* 81:75–98.
- Tlauka, M. 2004. Display-control compatibility: The relationship between performance and judgments of performance. *Ergonomics* 47:281–95.
- Tombu, M., and P. Jolicoeur. 2005. Testing the predictions of the central capacity sharing model. *J Exp Psychol Hum Percept Perform* 31:790–802.
- Treisman, A. M. 1964. Selective attention in man. *Br Med Bull* 20:12–6.
- Treisman, A. M., and G. Gelade. 1980. A feature-integration theory of attention. *Cogn Psychol* 12:97–136.
- Trimmel, M., and R. Huber. 1998. After-effects of human-computer interaction indicated by P300 of the event-related brain potential. *Ergonomics* 41:649–55.
- Tullis, T. S., F. J. Tranquada, and M. J. Siegel. 2011. Presentation of information. In *Handbook of Human Factors in Web Design*, 2nd ed., ed. K.-P. L. Vu and R. W. Proctor, 153–89. Boca Raton, FL: CRC Press.
- Usher, M., Z. Olami, and J. L. McClelland. 2002. Hick's law in a stochastic race model with speed-accuracy tradeoff. *J Math Psychol* 46:704–15.
- van de Laar, M. C., W. P. M. van den Wildenberg, G. J. M. van Boxtel, and M. W. van der Molen. 2010. Processing of global and selective stop signals application of Donders' subtraction method to stop-signal task performance. *Exp Psychol* 57:149–59.
- Van Zandt, T., H. Colonius, and R. W. Proctor. 2000. A comparison of two response time models applied to perceptual matching. *Psychon Bull Rev* 7:208–56.
- Verhaeghen, P., and A. Marcoen. 1996. On the mechanisms of plasticity in young and older adults after instruction in the method of loci: Evidence for an amplification model. *Psychol Aging* 11:164–78.
- Vu, K.-P. L., and R. W. Proctor. 2003. Naïve and experienced judgments of stimulus-response compatibility: Implications for interface design. *Ergonomics* 46:169–87.
- Vu, K.-P. L., and R. W. Proctor. 2004. Mixing compatible and incompatible mappings: Elimination, reduction, and enhancement of spatial compatibility effects. *Q J Exp Psychol* 57A:539–56.
- Vu, K.-P. L., and R. W. Proctor. 2008. Age differences in response selection for pure and mixed stimulus-response mappings and tasks. *Acta Psychol* 129:49–60.
- Vu, K.-P. L., R. W. Proctor, A. Bhargav-Spanzel, B. -L. Tai, J. Cook, and E. E. Schultz. 2007. Improving password security and memorability to protect personal and organizational information. *Int J Hum Comput Stud* 65:744–57.
- Walker, N., J. B. Smelcer, and E. Nilsen. 1991. Optimizing speed and accuracy of menu selection: A comparison of walking and pull-down menus. *Int J Man Mach Stud* 35:871–90.
- Wang, H., and R. W. Proctor. 1996. Stimulus-response compatibility as a function of stimulus code and response modality. *J Exp Psychol Hum Percept Perform* 22:1201–17.
- Ward, J. 2010. *The Student's Guide to Cognitive Neuroscience*. 2nd ed. New York: Psychology Press.
- Warm, J. S., R. Parasuraman, and G. Matthews. 2008. Vigilance requires hard mental work and is stressful. *Hum Factors* 50:433–41.
- Wickelgren, W. A. 1977. Speed-accuracy tradeoff and information processing dynamics. *Acta Psychol* 41:67–85.
- Wickens, D. D. 1970. Encoding categories of words: An empirical approach to meaning. *Psychol Rev* 77:1–15.
- Wickens, C. D. 1984. Processing resources in attention. In *Varieties of Attention*, ed. R. Parasuraman and D. R. Davies, 63–102. San Diego, CA: Academic Press.
- Wickens, C. D., S. Gordon, Y. Liu, and J. Lee. 2003. *An Introduction to Human Factors Engineering*. 2nd ed. New York: Prentice Hall.
- Wolfe, J. M. 2007. Guided Search 4.0: Current progress with a model of visual search. In *Integrated Models of Cognitive Systems*, ed. W. D. Gray, 99–119. New York: Oxford University Press.
- Wood, L. E., and J. D. Pratt. 1987. Pegword mnemonic as an aid to memory in the elderly: A comparison of four age groups. *Educ Gerontol* 13:325–39.
- Yantis, S. 2000. Goal-directed and stimulus-driven determinants of attentional control. In *Control of Cognitive Processes: Attention and Performance XVIII*, ed. S. Monsell and J. Driver, 195–208. Cambridge, MA: MIT Press.
- Young, P. M., B. A. Clegg, and C. A. P. Smith. 2004. Dynamic models of augmented cognition. *Int J Hum Comput Interact* 17:259–73.

3 Mental Models in Human–Computer Interaction

Stephen J. Payne

CONTENTS

3.1 What Is a Mental Model?.....	41
3.2 Mental Models of Text and Other Artifacts	45
3.3 Instructions for Mental Models.....	47
3.3.1 Multimedia Instruction	47
3.3.2 Theory of Learning by Not Doing	48
3.4 Shared Mental Models	50
References.....	52

The plan for this chapter is as follows. It begins by reviewing and discussing the term “mental models” as it has been used in the literature on human–computer interaction (HCI), and in the neighboring disciplines of cognitive psychology where it was first coined. There is little consensus on what exactly is and is not a mental model, and yet it is too widely used for any posthoc attempt at a narrower definition to somehow cleanse the field. In consequence, I characterize several layers of theoretical commitment that the term may embrace, following an earlier discussion (Payne 2003). To illustrate the argument, several classic and more recent studies from the HCI literature will be reviewed, with pointers to others. This first part of the chapter is based on material published in Payne (2003).

In cognitive psychology, mental models have major currency in two sub-disciplines—text comprehension and reasoning, although in the former they more often currently go by the name “situation models.” Discussion in the latter focuses on quite refined theoretical disputes that currently have little relevance for HCI. The work on text comprehension, however, is germane. With the advent of the web, the comprehension of text of various kinds has become a dominant mode of HCI, with important design issues for websites, digital libraries, and so on. Interaction with text is in some ways a paradigm for interaction with information. With these points in mind, the concept of mental models in text comprehension will be discussed, with a particular eye to the issues that HCI accentuates, such as understanding multiple texts.

Two of the major practical questions raised by mental models are (1) How are they acquired? and (2) How can their acquisition be supported by instruction? The third section of this chapter will discuss two angles on these questions in HCI: first, the use of interactive computation and multimedia as an instructional method; second, the important tension between exploration and instruction, first systematically

discussed in the HCI literature by Carroll’s (1990) work on minimalism.

Finally, the paper will review some recent work on the importance of mental models for understanding aspects of collaborative teamwork. This area suggests that a relatively expansive view of human knowledge representations may be necessary for progress in HCI.

Throughout the chapter, a particular approach is taken to review: to choose one or two key studies and report them in some detail. I hope that this will allow some of the empirical methodologies and the rich variation in these to be conveyed. The chosen studies will be accompanied by some further references to the literature, but there are too many subtopics reviewed to aim for completeness.

3.1 WHAT IS A MENTAL MODEL?

The user’s mental model of the device is one of the more widely discussed theoretical constructs in HCI. Alongside wide-ranging research literature, even commercial style guides have appealed to mental models for guidance (i.e., Mayhew 1992; Tognazzini 1992; Apple Human Interface Guidelines Apple Computer Inc. 1987).

Yet a casual inspection of the HCI literature reveals that mental models are used to label many different aspects of users’ knowledge about the systems they use. Nevertheless, I propose that even this simple core construct—what users know and believe about the systems they use—is worth highlighting and promoting. It is more distinctive than it might first seem, especially in comparison with other cognitive-science approaches. Further, beyond the core idea there is a progression of stronger theoretical commitments that have been mobilized by the mental models label, each of which speaks to important issues in HCI research, if not yet in practice.

The fundamental idea is that the *contents* of people's knowledge, including their theories and beliefs, can be an important explanatory concept for understanding users' behavior in relation to systems. This idea may seem obvious and straightforward, but in fact it suggests research questions that go against the grain of most contemporary cognitive psychology, which has concerned itself much more with the general limits of the human-information-processing system, such as the constraints on attention, retrieval, and processing. Thus, cognitive psychology tends to focus on the *structure* of the mind, rather than its contents. (The major exception to the rule that cognitive psychology has been obsessed with architecture over content is the work on expertise, and even here, recent work has focused on explanations of extreme performance in terms of general independent variables such as "motivated practice," i.e., Ericsson, Krampe, and Tesch-Romer [1993], rather than epistemological analysis.)

Refocusing attention on mental content about particular domains is what made mental models a popular idea in the early 1980s, such as the papers in Gentner and Stevens (1983). For example, work on naïve physics (i.e., McCloskey 1983) attempts to explain people's reasoning about the physical world, not in terms of working memory limits or particular representations, but in terms of their beliefs about the world, such as the nature of their theories of mechanics or electricity, for example. This focus on people's knowledge, theories, and beliefs about particular domains transfers naturally to questions in HCI, where practical interest may focus on how users conceive the workings of a particular device, how their beliefs shape their interactive behavior, and what lessons may be drawn for design.

In this mold, consider a very simple study of my own (Payne 1991). Students were interviewed about ATMs. Following Collins and Gentner (1987) among others, "what if" questions were posed to uncover student's theories about the design and function of ATMs. For example, students were asked whether machines sometimes took longer to process their interactions; what information was stored on the plastic card; and what would happen if they "typed ahead" without waiting for the next machine prompt.

The interviews uncovered a wide variety in students' beliefs about the design of ATMs. For example, some assumed that the plastic card was written to as well as read from during transactions, and thus could encode the current balance of their account. Others assumed that the only information on the card was the user's personal identification number, allowing the machine to check the identity of the user (as it turns out, both these beliefs are incorrect). A conclusion from this simple observation is that users of machines are eager to form explanatory models and will readily go beyond available data to infer models that are consistent with their experiences. (One might wonder whether such explanations were not merely ad hoc, prompted during the interview: in fact some were, but explicit linguistic cues—such as "I've always thought"—strongly suggested that many were not.)

Another observation concerning students' "models" of ATMs was that they were fragmentary, perhaps more

fragmentary than the term "model" might ordinarily connote: they were collections of beliefs about parts of the system, processes, or behaviors, rather than unified models of the whole design. Students would happily recruit an analogy to explain one part of the machine's operation that bore no relation to the rest of the system. This fragmentary character of mental models of complex systems may be an important aspect (see i.e., Norman 1983), allowing partial understandings to be maintained. One implication is that users' mental models of single processes or operations might be a worthwhile topic for study and practical intervention (in design or instruction).

One widely held belief about a particular process affected the students' behavior as users. Almost all respondents believed that it was not possible to type ahead during machine pauses. At the time the study was conducted this was true for some, but not all, designs in use. Consequently, in some cases transactions were presumably being needlessly slowed because of an aspect of users' mental models.

A more recent study of a similar kind is an investigation of users' models of the navigation facilities provided by Internet browsers (Cockburn and Jones 1996). Internet browsers, like Internet Explorer, maintain history lists of recently visited pages, providing direct access to these pages without needing to enter the URL or follow a hyperlink. The "back" and "forward" buttons provide a very frequently used mechanism for browsing history lists, but do users have good mental models for how they work? Cockburn and Jones (1996) showed that many do not.

The history list of visited pages can be thought of as a stack: a simple last-in-first-out data structure to which elements can be added (pushed) or taken out (popped) only from the top (consider a stack of trays in a canteen). When a new web page is visited by following a hyperlink, or by entering a URL, its address is pushed onto the top of the stack. This is true even if the page is already in the history list, so that the history list may contain more than one copy of the same page. However, when a page is visited by using the Back button (or, at least typically, by choosing from the history list), the page is not pushed onto the stack. So, what happens when the currently displayed page is not at the top of the stack (because it has been visited via the history list) and a new link is followed (or a new URL entered)? The answer is that all the pages in the history list that were above the current page are popped from the stack, and the newly visited page is pushed onto the stack in their place. For this reason the history list does *not* represent a complete record, or time-line of visited pages, and not all pages in the current browsing episode can be backed-up to. In Cockburn and Jones' study, few users appreciated this aspect of the device.

This then, has been the major thrust of work on mental models in HCI: what do people know and believe to be true about the way the systems they interact with are structured? How do their beliefs affect their behavior? In this literature a "mental model" is little more than a pointer to the relevant parts of the user's knowledge, yet this is not to deny its usefulness. One approach that it has engendered is a typology of

knowledge—making groupings and distinctions about types of knowledge that are relevant in certain circumstances. It is in exactly this way that a literature on “shared mental models” as an explanatory concept in teamwork has been developed. This topic is perhaps the most rapidly growing area of mental models research in HCI and will be reviewed in the final section of this chapter.

However, as argued at length in Payne (2003), there are approaches to mental models in HCI that go beyond a concern with user knowledge and beliefs to ask more nuanced theoretical questions. The first of these is to investigate the form of mental models by inspecting the processes through which mental models might have their effects on behavior.

A powerful idea here is that mental models of machines provide a problem space that allows more elaborate encoding of remembered methods, and in which novice or expert problem solvers can search for new methods to achieve tasks.

The classic example of this approach is the work of Halasz and Moran (1983) on Reverse Polish Notation (RPN) calculators. RPN is a post-fix notation for arithmetic, so that to express $3 + 4$, one would write $3\ 4\ +$. RPN does away with the need for parentheses to disambiguate composed operations. For example $(1 + 2) * 3$ can be expressed $1\ 2\ +\ 3\ *$ with no ambiguity. RPN calculators need a key to act as a separator between operands, which is conventionally labeled ENTER, but they do not need an = key, as the current total can be computed and displayed whenever an operator is entered.

Halasz and Moran taught one group of students how to use an RPN calculator using instructions, like a more elaborate version of the introduction above, which simply described the appropriate syntax for arithmetic expressions. A second group of subjects was instructed, using a diagram, about the stack model that underlies RPN calculation. Briefly, when a number is keyed in, it is “pushed” on top of a stack-data structure (and the top slot is displayed). The ENTER key copies the contents of the top slot down to the next slot. Any binary arithmetic operation is always performed on the contents of the top two slots and leads to the result being in the top slot, with the contents of slots 3 and below moving up the stack.

Halasz and Moran discovered that the stack-model instructions made no difference to participants’ ability to solve routine arithmetic tasks: the syntactic “method-based” instructions sufficed to allow participants to transform the tasks into RPN notation. However, for more creative problems (such as calculating $(6 + 4)$ and $(6 + 3)$ and $(6 + 2)$ and only keying the number 6 once) the stack group was substantially better. Verbal protocols showed that these subjects reasoned about such problems by mentally stepping through the transformations to the stack at each keystroke.

This kind of reasoning, stepping through a sequence of states in some mental model of a machine, is often called “mental simulation” in the mental models literature, and the kind of model that allows simulation is often called a “surrogate” (Young 1983; Carroll and Olson 1988). From a practical standpoint, the key property of this kind of reasoning is that it results in behavior that is richer and more flexible than the mere rote following of learned methods. The idea that the

same method may be encoded more richly, so that it is more flexible and less prone to forgetting will be returned to later in the chapter when a theory of mental models of interactive artifacts is considered, and when ideas about instruction for mental models are reviewed.

A second example of mental models providing a problem space elaboration of rote methods comes in the work of Kieras and Bovair (1984). This research was similar to that of Halasz and Moran (1983) in that it compared the learning performance of two groups: (1) one instructed with rote procedures, (2) the other additionally with a diagrammatic model of the device on which the procedures were enacted. In this case, the device was a simple control panel, in which each rote procedure specified a sequence of button-pushes and knob-positions leading to a sequence of light-illuminations. The model was a circuit diagram showing the connections between power-source switches and display-lights.

Kieras and Bovair (1984) found that the participants instructed with the model learned the procedures faster, retained the procedures more accurately, executed the procedures faster, and could simplify inefficient procedures that contained redundant switch settings. They argued that this was because the model (circuit diagram) explained the contingencies in the rote-action sequences (i.e., if a switch is set to MA, so that the main accumulator circuit is selected, then the FM, fire main, button must be used).

A related theoretical idea is that mental models are a special kind of representation, sometimes called an *analog* representation: one that shares the structure of the world it represents. This was taken as the definitional property of mental models by the modern originator of the term, the British psychologist Kenneth Craik (1943). It is this intuition that encourages the use of terms like “mental simulation”—the intuition that a mental model is like a physical model, approximating the structure of what it represents, just as a model train incorporates (aspects of) the physical structure of a train.

The idea that mental models are analog in this sense is a definitional property in the work on reasoning and comprehension by Johnson-Laird (Johnson-Laird 1983, 1989; this will be further discussed in Section 3.2, concerning representations of text) and also in the theory of Holland et al. (1986) and Moray (1999). However, there are different nuances to the claim, which must be considered. And, in addition, there is a vexed question to be asked; namely, what is the explanatory or predictive force of a commitment to analog representational form? Is there any reason for HCI researchers to pay attention to theoretical questions at this level?

Certainly, this is the view of Moray (1999) who is concerned with mental models of complex dynamic systems, such as industrial plants. He proposes that models of such systems are structure-sharing *homomorphisms* rather than isomorphisms, that is, they are many to one rather than one-to-one mappings of objects, properties, and relations. (In this he follows Holland et al. 1986.)

Homomorphic models of dynamic systems may not share structure with the system at the level of static relations, but

only at the level of state-changes. Thus, such models have the character of state-transition diagrams, making the empirical consequences of structure sharing somewhat unclear, because any problem space can be represented in this way.

In my view, a clearer view of the explanatory force of analog mental models can be derived by carefully considering the ideas of computational and informational equivalence first introduced by Simon (1978).

It is obviously possible to have two or more distinct representations of the same information. Call such representations “informationally equivalent” if all the information in one is inferable from the other, and vice versa. Two informationally equivalent representations may or may not additionally be “computationally equivalent,” meaning that the cost structure of accessing and processing the information is equivalent in both cases, or, as Larkin and Simon (1987) put it: “information given explicitly in the one can also be drawn easily and quickly from the information given explicitly in the other, and vice versa.” As Larkin and Simon point out, “easily” and “quickly” are not precise terms, and so this definition of computational equivalence is inherently somewhat vague; nevertheless it points to empirical consequences of a representation (together with the processes that operate upon it) that depend on form, and therefore go beyond mere informational content.

In Payne (2003), I propose adopting *task-relative* versions of the concepts of informational and computational equivalence. Thus, representations are informationally equivalent, *with respect to a set of tasks*, if they allow the same tasks to be performed (i.e. contain the requisite information for those tasks). The representations are, additionally, computationally equivalent with respect to the tasks they allow to be performed, if the *relative difficulty* of the tasks is the same, whichever representation is being used. (Note that according to these definitions, two representations might be computationally equivalent with regard to a subset of the tasks they support but not with regard to the total set, so that in Larkin and Simon’s sense they would merely be informationally equivalent. The task-relative versions of the constructs thus allow more finely graded comparisons between representations.)

This idea can express what is behaviorally important about the idea of analog models, or structure-sharing mental representations of a state of affairs of a dynamic system. An analog representation is computationally equivalent (with respect to some tasks) to external perception and manipulation of the state of affairs it represents.

Bibby and Payne (1993, 1996) exploited this distinction between computational and informational equivalence in the domain of HCI, using a computer simulation of a device derived from that studied by Kieras and Bovair (1984). The device was a multiroute circuit, in which setting switches into one of several configurations would make a laser fire; various indicator lights showed which components of the circuit were receiving power. What concerned Bibby and Payne (1993) was the idea of computational equivalence between a mental model and a diagram of the device, rather than the device itself.

Bibby and Payne asked participants to repeatedly perform two types of tasks: a switch task, in which all but one switch was already in position to make a laser fire (the participant had to key the final switch) and a fault task, in which the pattern of indicator lights was such that one of the components must be broken (the participant had to key the name of the broken component).

Participants were instructed about the device with either a table, which showed the conditions under which each indicator light would be illuminated, or with procedures, sequences of switch positions enabling the laser to be fired. Both instructions were sufficient for both switch and fault tasks; they were informationally equivalent with respect to those tasks. However, the table made the fault task easier than the switch task, whereas the procedures made the switch task easier.

During practice, when participants consulted the instructions, this pattern of relative difficulty was confirmed by a crossover interaction in response times. Furthermore, when the instructions were removed from the participants, so that they had to rely on their mental representation of the device, the crossover interaction persevered, demonstrating that the mental representations were computationally equivalent to the external instructions.

In subsequent experiments, Bibby and Payne (1996) demonstrated that this pattern persevered even after considerable interaction with the device that might have been expected to provide an opportunity to overcome the representational constraints of the initial instruction. The crossover interaction eventually disappeared only after extended practice on the particular fault-and-switch task (80 examples of each: perhaps because of asymptotic performance having been reached). At this point, Bibby and Payne introduced two similar but new types of tasks designed so that once again, the table favored one task whereas procedures favored the other. (However, the device instructions were *not* re-presented.) At this point the crossover re-appeared, demonstrating that participants were consulting their instructionally derived mental model of the device, and that this was still in a form computationally equivalent to the original external representation of the instructions.

Practically, this research shows that the exact form of instructions may exert long-lasting effects on the strategies that are used to perform tasks, so that designers of such instructions must be sensitive not only to their informational content but also to their computational properties. In this light, they also suggest that one instructional representation of a device is very unlikely to be an optimal vehicle for supporting all user tasks: it may well be better to provide different representations of the same information, each tailored to particular tasks. In this sense, perhaps instructions should mirror and exploit the natural tendency, noted above, for users to form fragmentary mental models, with different fragments for different purposes.

In terms of theory, Bibby and Payne’s findings lend support to the suggestion developed above that mental models of a device that are formed from instructions may be computationally equivalent to the external representations of the

device. This idea gives a rather new reading, and one with more ready empirical consequences to the theoretically strong position that mental models are essentially analog, homomorphic representations.

3.2 MENTAL MODELS OF TEXT AND OTHER ARTIFACTS

The psychological literature on text comprehension has been transformed by the idea of a situation model, first put forward as part of a general theory of text comprehension by van Dijk and Kintsch (1983), and developed over the years by Kintsch (1998) and followers. The central idea of the general theory is that readers construct mental representations of what they read at several different levels. First, they encode the surface form of the text: the words and syntax. Second, they go beyond this to a representation of the propositional content of the text. Finally, they go beyond the propositional context of the text itself to represent what the text is about, incorporating their world knowledge to construct a situation model or mental model of the described situation.

(Under this view, it is the content that distinguishes a situation model from a text base, rather than a representational format. However, some researchers, notably Johnson-Laird (1983), and followers have pursued the idea of mental models derived from text as analog representations of the described situation. Thus, in text comprehension, there is a version of the issue discussed in part one.)

It is instructive to consider some of the evidence for situation models, and what important issues in text comprehension the theory of situation models allows us to address.

A classic early study was conducted by Bransford, Barclay, and Franks (1972). They asked participants to read simple sentences such as,

Three turtles rested beside/on a floating log, and a fish swam beneath them.

(The slash indicates that some subjects read the sentence with the word “beside,” and others read the same sentence with the word “on.”)

In a later recognition test, interest centered on how likely readers were to falsely accept minor rewordings of the original sentences. In the above case, the foil sentence was

Three turtles rested beside/on a floating log, and a fish swam beneath it.

The key finding was that people who had read the “on” versions of the sentences were much more likely to accept the changed version of the sentence, despite the fact that at the level of the sentences the difference between original and foil sentences in the two conditions is identical, limited in each case to the last word of the sentence. The reason for false recognition in one case is because, in this case, but not when “on” is replaced by “beside,” the original and foil sentences describe the same situation.

A related series of experiments was reported by Fletcher and Chrysler (1990). In a series of carefully controlled

experiments, they varied the overlap between sentences in a recognition test and sentences from 10 different texts read by the participants. Each text described a state of affairs (i.e., the relative cost of antiques) consistent with a linear ordering among a set of five objects. They found that participants were influenced by overlap between sentences at study and test corresponding to the three levels of discourse representation proposed by van Dijk and Kintsch (1983): surface form, text base, and situation model. Recognition performance was best when distracter items were inconsistent with all three levels of representation. Recognition was above chance when distracters violated merely the surface form of the original sentences (i.e. substituting rug for carpet). It improved further when propositional information from the text base, but not the linear ordering of the situation, was violated. Recognition was best of all when the distracters were inconsistent with the situation described by the text. This suggests that some aspects of the structure of the situation (in this case a set of linear orderings) were retained.

Next, consider work by Radvansky and Zacks (Radvansky and Zacks 1991; Radvansky, Spieler, and Zacks 1993). In these experiments, participants read sentences such as, “The cola machine is in the hotel,” each of which specified the location of an object. In one condition sentences shared a common object (i.e. cola machine) but different locations. In a second condition, different objects share a common location (i.e. the city hall). Later in the experiment participants were given a speeded-recognition test. Radvansky and Zacks found a significant fan effect for the common object condition; times to verify sentences increased as the number of different locations rose. For the common location sentences no significant fan effect emerged. This was interpreted as evidence that participants formed mental models around the common location (a representation of such a location containing all the specified objects) and retrieval from long-term memory (LTM) was organized around these mental models. It is impossible, or much harder, to form such a representation of the same object in multiple locations.

What all these studies, and many like them, reveal is that when understanding text, readers spontaneously construct a mental representation that goes beyond the text itself and what it means, and use inferences to construct a richer model of what the text is about—a situation model.

Beyond these refined and clever, but undeniably rather narrow experimental contexts, the construct of situation models has been put to work to illuminate some practical issues concerning text comprehension, and exactly this issue will be returned to later, where we will see how it can inform attempts to understand instructional strategies for engendering useful mental models.

There are two principal ways in which the literature on text comprehension is relevant to HCI. First, it provides support for the idea that a mental model is a representation of what a representational artifact represents. The layered model of text comprehension previously outlined can be generalized to the claim that the user of any representational artifact must construct a representation of the artifact itself, and of what

the artifact represents, and of the mapping between the two (how the artifact represents). This is the basis of the Yoked State Space (YSS) hypothesis (Payne, Squibb, and Howes 1990).

If a reader's goal is just to understand a text, as it was in the experiments just reviewed, then the text-representation can be discarded once a model has been constructed. However, there are many tasks of text *use*, in which it is necessary to maintain a representation of the text, alongside a mental model of the meaning of the text. Consider, for example, the tasks of writing and editing, or of searching for particular content in a text. In such tasks, it is necessary to keep in mind the relation between the surface form of the text—wording, spatial layout, and so on—and its meaning. Text is a representational artifact, and to *use* it in this sense one needs a mental representation of the structure of the text, and of the “situation” described by the text and of the mapping between the two.

According to the YSS hypothesis (Payne, Squibb, and Howes 1990), this requirement is general to all representational artifacts, including computer systems. To use such artifacts requires some representation of the domain of application of the artifact—the concepts the artifact allows you to represent and process. The user's goals are states in this domain, which is therefore called the goal space. However, states in the goal space cannot be manipulated directly. Instead, the user interacts with the artifact, and therefore needs knowledge of the artifact, and of the operations that allow states of the artifact to be transformed. Call this problem space the device space. In order to solve problems in the goal space by searching in the device space, the user must know how the device space represents the goal space. In this sense the two spaces need to be yoked. The minimal device space for a certain set of tasks must be capable of representing all the states in the corresponding goal space. More elaborate device spaces may incorporate device states that do not directly represent goal states, but which allow more efficient performance of tasks, just as the stack model of an RPN calculator allows an elaboration of methods for simple arithmetic.

The work of Halasz and Moran (1983) can readily be assimilated into the YSS framework. The no-model condition was provided with enough information to translate algebraic expressions into their Reverse Polish equivalent. However, in this understanding of RP expressions, the ENTER key was given merely an operational account, serving simply as a separator of operands, and did not transform the device state. The stack model, however, provides a figurative account of the ENTER key.

This discussion illustrates a practical lesson for the design of interfaces and instructions. In the case of the copy buffer and the calculator stack, the standard interface does not allow the appropriate device space readily to be induced, so that conceptual instructions must fill the gap. The obvious alternative, which has been developed to some extent in both cases, is to redesign the user interface so as to make the appropriate device space visible. These examples suggest a simple heuristic for the provision

of conceptual instructions that may help overcome the considerable controversy over whether or not such instructions (as opposed to simple procedural instructions) are useful. According to this heuristic, conceptual instructions will be useful if they support construction of a YSS that the user would otherwise have difficulty inducing (Payne, Howes, and Hill 1992).

A more direct way in which text comprehension research is relevant to HCI is that so much HCI is reading text. Beyond the standard issues, the widespread availability of electronic texts raises some new concerns that have not yet seen much work, yet are perhaps the most directly relevant to HCI design. Two issues stand out: (1) the usability of documents that incorporate multiple media alongside text, and (2) the exploitation by readers of multiple texts on the same topic.

How are multimedia “texts” that incorporate graphics comprehended? There is only a small literature on this within the mainstream field of text comprehension, but this literature exploits the idea of a mental model.

Glenberg and Langston (1992) argued that the widespread idea that diagrams can assist the comprehension of technical text had, at the time, been little tested or understood and that mental models were an important explanatory construct. In their analysis, diagrams are useful in concert with texts precisely because they assist the construction of mental models. This idea has been pursued in a very active program of work on multimedia instruction by Mayer and colleagues, which will be reviewed in the next section.

What about when the multiple sources of information are not presented as part of a single text, but rather independently, covering overlapping ground, so that the reader has to perform all the integration and mapping? This is the issue of multiple texts, and it has become commonplace in the age of the Internet. It is now rarely the case that a student struggles to find relevant source documents on a topic. Instead, students are typically faced with an overabundance of relevant materials and must somehow allocate their time across them, and integrate the knowledge they derive from different sources.

Perfetti (1997) has suggested that learning from multiple texts is one of the most important new challenges for text researchers. Research has shown, for example, that integrating information across multiple texts is a skill that does not come readily but can be acquired and taught (Stahl et al. 1996; Rouet et al. 1997).

The YSS theory raises important issues here. As previously noted, everyday reading of text can be seen as engendering a progression of mental representations moving from the surface form through the propositional content to a situation model. When reading, earlier representations can be discarded as later ones are formed, but for other tasks of text use, the reader needs to maintain a representation of the form of the multitext, and map this form onto the content. Payne and Reader (in press) refer to such a representation as a structure map.

The usefulness of a structure map becomes even more apparent when multiple texts are considered. In this case, structure maps could play a role in encoding *source* information, which

might be important not only for locating information, but also for integrating diverse and potentially contradictory information and for making judgments of trust or confidence in the information. Source information might additionally encode temporal properties of information sources, and thus be useful for memory updating—revising knowledge in the light of new information, making distinctions between current and super-seeded propositions.

The widespread availability of the web not only means that multiple texts are more widely encountered, but also encourages a situation where multiple texts are read in an interleaved fashion, in a single sitting, or at least temporally close, raising the importance of the above challenges, and meaning that recency in autobiographical memory is unlikely to accomplish source identification, so further stressing the importance of a structure map.

Payne and Reader (in press) studied readers' ability to search for specific ideas in multiple texts that they had just read. They found evidence that readers spontaneously constructed structure maps, as just described, in that they showed some memory of which documents contained which ideas, even when they did not expect to need such knowledge when reading the texts.

3.3 INSTRUCTIONS FOR MENTAL MODELS

3.3.1 MULTIMEDIA INSTRUCTION

If mental models are important for operating devices, how should they best be taught? We have seen that models are constructed automatically by readers of text, but can modern computational media, such as animations, be used to improve the acquisition of mental models from instructional texts, just as Glenberg and Langston (1992) suggested in the case of simple diagrams? Just such a question has been addressed in a long-standing program of work by Richard Mayer and colleagues, which will be reviewed in this section.

Mayer and Moreno (2002) present a cognitive theory of multimedia learning, which builds on three main ideas:

1. From dual coding theory the authors suppose that humans have separate visual and verbal information processing systems (Clark and Paivio 1991; Paivio 1986)
2. From cognitive load theory the authors assume that the processing capacity of both the visual and the verbal memory system is strictly limited (Baddeley 1992; Chandler and Sweller 1991) and that cognitive load during instruction can interfere with learning
3. From constructivist learning theory the authors take the idea that meaningful learning requires learners actively to select relevant information, to structure it into coherent representations, and make connections with other relevant knowledge (Mayer 1996, 1999a)

This latter process, of building coherent representations that connect information from different modalities with

pre-existing knowledge, bears clear relation to Johnson-Laird's construct of mental models, and indeed Mayer and colleagues use the term in this context. In the case of the physical systems that many of their studies have addressed, mental models may take the form of cause-effect chains. According to Mayer and Moreno (2002) a key design principle for instructional materials is that they should maximize the opportunity for these model-construction processes to be completed.

Mayer and colleagues have conducted a large number of experiments comparing learning from multimedia source materials with learning from components of these materials (words, pictures, etc.) successively or in other kinds of combination. Based on this research, Mayer (1999b) and Mayer and Moreno (2002) have identified some principles of instructional design that foster multimedia learning.

The *multiple presentation principle* states that explanations in words and pictures will be more effective than explanations that use only words (Mayer and Moreno 2002, p. 107). When words only are presented, learners may find it difficult to construct an appropriate mental image, and this difficulty may block effective learning. Mayer and Anderson (1991; Experiment 2b) compared four treatment groups: (1) words with pictures, (2) words only, (3) pictures only, and (4) control, on tests of creative problem solving involving reasoning how a bicycle pump works. Results demonstrated that participants in the words with pictures group generated a greater number of creative problem solutions than did participants in the other groups. Interestingly, animation without narration was equivalent to no instruction at all. Other studies have offered support for the general idea that learners will acquire richer knowledge from narration and animation than from narration alone (Mayer and Anderson 1991, Experiment 2a; Mayer and Anderson 1992, Experiments 1 and 2).

The *contiguity principle* is the claim that simultaneous as opposed to successive presentation of visual and verbal materials is preferred (Mayer and Moreno 2002), because this will enable learners to build referential connections more readily (Mayer and Sims 1994). Mayer and Anderson (1991, Experiments 1 and 2) studied a computer-based animation of how a bicycle pump works. They compared a version that presented words with pictures against the same content presenting words before pictures, and tested acquisition with tests of creative problem solving. Those in the words-with-pictures group generated about 50% more solutions to the test problems than did subjects in the words-before-pictures group.

The *individual differences principle* predicts that factors such as prior knowledge or spatial ability will influence transfer of learning from multimedia materials, moderating the effects of other principles (Mayer 1999c). With regard to domain specific knowledge, Mayer proposed that experienced learners may suffer little decrease in problem solving transfer when receiving narration and animation successively because their background knowledge will allow a mental model to be constructed from the words alone, then linked to the visual information. Low-experience learners, on the

other hand, will have no means to over-ride the effects underlying the contiguity principle, and their problem solving transfer will suffer (Mayer and Sims 1994). In support of this suggestion, experimental work by Mayer and Gallini (1990) demonstrated across three studies that the synchronization of words and pictures served to improve transfer for low- but not high-experience learners.

The *chunking principle* refers to a situation in which visual and verbal information must be presented successively, or alternately (against the contiguity principle). It states that learners will demonstrate better learning when such alternation takes place in short rather than long segments. The reasoning is straightforward, given the assumptions of the framework: working memory may become overloaded by having to hold large chunks before connections can be formed (Mayer 1999b). An experiment by Mayer and Moreno (1998) investigated this chunking principle using explanations of how lightning storms develop. The ability to solve novel, transfer problems about lightning exhibited by a ‘large chunk’ group (who received all the visual information before or after all the verbal information) was compared with that of a ‘small chunk’ group (alternating presentations of a short portion of visual followed by a short portion of narration). The gain in performance of the small chunk group over the large chunk group was circa 100% (Mayer and Moreno 1998).

The debt of Mayer’s work to Sweller’s program of research on Cognitive Load Theory is obvious. Mayer’s design principles reflect the premise that students will learn more deeply when their visual and/or verbal memories are not overloaded. Students are better able to make sense of information when they receive both verbal and visual representations rather than only verbal; when they can hold relevant visual and verbal representations in working memory at the same time; when they have domain specific knowledge and/or high spatial ability; and when they receive small bits of information at a time from each mode of presentation.

Despite incredibly positive research results, at this stage Mayer’s work should be viewed with a little caution. Almost all of the experiments utilize very short instructional presentations, with some of the animations lasting only 30 seconds. Subjects are then required to answer problem-solving questions that seem ambiguous, requiring students to be fairly creative in order to generate solutions. Mayer’s work also typically neglects to include any tests of long-term retention. It may conceivably be falling into the instructional trap of maximizing performance during learning at the expense of longer-term performance. This issue is the focus of the next section.

3.3.2 THEORY OF LEARNING BY NOT DOING

Mayer’s theory of multimedia instruction adheres to the common assumption that the optimal design of instructional material involves minimizing the cognitive burden on the learner due to the limits of the working memory.

Yet minimizing the mental effort of learners is not necessarily or always a good instructional strategy. According

to Schmidt and Bjork (1992), instructional conditions that achieve the training goals of generalizability and long-term retention are not necessarily those that maximize performance during the acquisition phase.

They argue that the goal of instruction and training in real-world settings should first be to support a level of performance in the long term, and second to support the capability to transfer that training to novel-tasks environments. Methodologically, in order to measure a genuine *learning effect*, some form of long-term assessment of retention must take place; skill acquisition is not a reliable indicator of learning.

Schmidt and Bjork (1992) discussed three situations in which introducing difficulties for the learner can enhance long-term learning. First, studies that vary the scheduling of tasks during practice were reported. Random practice is more difficult than blocked schedules of practice, as a given task is never practiced on the successive trial. Using a complex motor task involving picking up a tennis ball and using it to knock over a particular set of barriers, Shea and Morgan (1979) reported a clear advantage for subjects who practiced under blocked conditions (subsets of barriers to knock), in terms of performance during practice. However, the amount of learning as demonstrated by the retention phase favored the random condition. Similar results have been reported by Baddeley and Longman (1978), Lee and Magill (1983), and (with verbal tasks) Landauer and Bjork (1978).

Schmidt and Bjork offer an explanation for this paradigm, in which retrieval practice may play a key role. They suggest that there may be a benefit, in terms of long-term retention, for activities that actually cause forgetting of the information to be recalled, forcing the learner to practice retrieving this information (Bjork and Allen 1970).

Experiments that vary the feedback the learner receives have demonstrated a similar phenomenon. A study by Schmidt et al. (1989) demonstrated that delaying the feedback that subjects received during motor tasks interfered with performance. However, on a delayed-retention test, those who had received the feedback least often demonstrated the most effective performance. This seems to contradict the established opinion that feedback is vital for effective learning. Schmidt and Bjork (1992) suggested that frequent feedback may actually serve to block information-processing activities that are important during the skill-acquisition phase.

A final area reviewed by Schmidt and Bjork concerns the introduction of variability during practice, such as when practicing tossing a beanbag at a target at a particular distance. Practicing at variable distances is more effective than practicing at a fixed distance (Kerr and Booth 1978).

Does the Schmidt and Bjork approach extend to HCI tasks, and in particular to instruction for mental models?

One impressive example of an instructional effect in the Schmidt and Bjork (1992) paradigm is informed by the idea of mental models or situation models derived from text, as discussed in Section 3.2 of this chapter. Informed by the distinction between a text base and a situation model, work by McNamara et al. (1996) has shown how expository text can

be designed to introduce difficulties for readers in exactly the productive manner advocated by the Schmidt and Bjork conception of training. These authors created two versions of target texts, one more coherent than the other (one experiment used a text about traits of mammals, a second used a text about heart disease). Coherence cues were provided by linking clauses with appropriate connectives and by inserting topic headings. The level of readers' background knowledge on the topic of the text was also assessed with a pretest. After reading a text, participants were given tests of the text base (free recall of the text propositions and specific factual questions about the contents of the text) and tests of the situation model (problem-solving-based questions, questions requiring inferences from the text, and a concept-sorting task).

McNamara et al. (1996) reported that for measures that tested the text base, the high coherence texts produced better performance. However, for situation-model measures, test performance for high-knowledge readers was better when they read the low-coherence text. McNamara et al. argued that limiting the coherence of a text forced readers to engage in compensatory processing to infer unstated relations in the text. This compensatory processing supported a deeper understanding of the text, in that the information in the text became more integrated with background knowledge. Thus, for high-knowledge readers, the texts that were more difficult to read improved the situation model by encouraging more transfer-appropriate processing. Low-knowledge readers were, presumably, unable to achieve the compensatory inferences, and therefore did better with more coherent texts. Because the text base does not incorporate background knowledge, it was not enhanced by any compensatory processing. (This finding is related to the work of Mayer and Sims [1994] reviewed above.)

One very successful practical approach to the design of instructions for interactive devices which is well known in the HCI community, is perhaps quite strongly related to this more theoretically oriented work. The concept of a "minimal manual" was outlined by Carroll (1990). It sought to minimize the extent to which instructional materials obstruct learning. Crucially, a well-designed Minimal Manual does not necessarily optimize the speed at which users can perform procedures as they read. Carroll's manuals avoided explicit descriptions that encouraged rapid but mindless rote performance. Instead, the emphasis was on active learning whereby learners were encouraged to generate their own solutions to meaningful tasks. This process was facilitated in part by reducing the amount of text provided and including information about error recovery.

O'Hara and Payne (1998, 1999) argued that learning from a problem-solving experience might be enhanced to the extent that problem solvers planned their moves through the problem space. Many puzzles with an interactive user interface, and indeed many user interfaces to commercial systems, encourage a one-step-at-a-time approach to problem solving, in which a move is chosen from the currently available set. This may be quick and relatively effortless, yet lead to little learning and inefficient solutions. For example, in an

HCI task, participants had to copy names and addresses from a database to a set of letters. Each item category from the database had to be copied to several letters, so that the most obvious and perhaps least effortful strategy of preparing letters one at a time was inefficient in terms of database access. O'Hara and Payne's manipulation was to increase the cost of making each move (in the copying experiment by adding a system lock-out time). This resulted in more planning, more think-time per move, meaning slower solutions in the first instance, but more efficient behavior in the long term, and the discovery of strategies that required fewer database accesses and fewer user inputs.

Recent work by Duggan and Payne (2001) combined several of the insights in the work just reviewed to explore acquisition of interactive procedures during instruction following. Good procedural instructions for interactive devices must satisfy two criteria. First, they must support performance. Like all procedural instructions they should effectively communicate the procedure they describe, so as to allow users who don't know the procedure to enact it successfully and efficiently. Second, they must support learning. In common with instructions for all procedures that will be used repeatedly, they should facilitate subsequent memory for the procedure, so that it might later be performed without consulting the instructions.

How might procedural instructions be designed so as to follow the Schmidt and Bjork paradigm and provide transfer-appropriate practice opportunities for the learner? Of course, not all manipulations that introduce difficulties during learning are beneficial for the learner. Simply making the instructions unclear is unlikely to be effective. However, much this idea may have informed the design of some commercial user manuals. The criterion that quality instructions must communicate the procedure that they describe cannot be ignored.

The work of Diehl and Mills (1995) further illustrated the relevance of the theory of text comprehension to the design of instruction for interactive procedures. They argued that in the case of procedural instructions the distinction between situation model and text base maps directly onto a distinction between memory for the procedure (as tested by later task performance) and memory for the instructions themselves.

Texts describing how to complete a task using a device (setting an alarm clock or constructing a child's toy) were provided. While reading a text, participants were required to either perform the task (read and do), or do nothing (read only). (In addition, Diehl and Mills studied some intermediate conditions, such as read and watch experimenter. These conditions produced intermediate results and are not relevant to the current argument.) The effect of these training methods was then examined by asking participants to recall the text and then complete the task.

Diehl and Mills reported that the increased exposure to the device in the read-and-do condition resulted in improved task performance times relative to the read-only condition. However, text recall was better in the read-only condition, supporting the conceptual separation of text base and situation model.

Inspired by this work, Duggan and Payne (2001) introduced a particular technique to exploit the principle of Schmidt and Bjork (1992) and the methods of McNamara and colleagues (1996). Like the manipulations of Diehl and Mills (1995), their innovation centered not on the design of the instructions per se, but rather on the way the instructions are read and used. Diehl and Mills' reported advantage for reading and doing over reading alone has no real practical implication, as it is difficult to imagine anyone advocating isolated reading as a preferred method. However, Duggan and Payne suggested that the way learners manage the interleaving of reading and doing will affect their later retention, and thus offers an important lever for improving instruction.

Many procedural instructions have a natural step-wise structure, and in these cases it is possible to execute the procedure while reading with minimal load on memory. Learners can read a single step, and then execute it before reading the next step. Such an approach is low in effort (and therefore attractive to the learner), but also low in transfer-appropriate practice and therefore, one would argue on the basis of the reviewed work, poor at encouraging retention. If learners could instead be prompted to read several procedural steps before enacting them, performance would be made more effortful, but learning might benefit. Readers would be encouraged to integrate the information across the chunk of procedural steps, and the increased memory load would provide transfer-appropriate practice.

Duggan and Payne (2001) developed this idea as follows. First, by implementing an online help system in the context of experimental tasks (programming a VCR) they forced participants into either a step-wise or a chunk-based strategy for interleaving reading and acting. These experiments demonstrated that reading by chunks did tax performance during training, but improved learning, in particular retention of the procedure. Next, they developed a more subtle, indirect manipulation of chunking. By adding a simple cost to the access of online instructions (c.f., O'Hara and Payne 1998), they encouraged readers to chunk steps so as to minimize the number of times the instructions were accessed. Just as with enforced chunking, this led to improved retention of the procedures.

3.4 SHARED MENTAL MODELS

In the last 10 years or so there has been a rapid surge of interest in the concept of shared mental models in the domain of teamwork and collaboration. The use of mental models in this literature, to date, is somewhat inexact, with little theoretical force, except to denote a concern with what the team members know, believe, and want. As the name suggests, shared mental models refers to the overlap in individuals' knowledge and beliefs.

The central thesis and motive force of the literature is that team performance will improve when team members share relevant knowledge and beliefs about their situation, task, equipment, and team. Different investigations and different

authors have stressed different aspects of knowledge, and indeed proposed different partitions into knowledge domains. (And recently, as we shall see, some investigators have questioned the extent to which overlapping knowledge is a good thing. There are some situations in which explicit distribution or division of knowledge may serve the team goals better.)

At first glance, the idea that teams need to agree about or share important knowledge seems intuitively plain. Models of communication (i.e., Clark 1992) stress the construction of a common ground of assumptions about each partner's background and intentions. The idea of shared mental models develops this idea in a plausible practical direction.

A recent study by Mathieu et al. (2000) was one of the most compelling demonstrations of the basic phenomenon under investigation, as well as being centered on an HCI paradigm. For these reasons, this study will be described and used as a framework to introduce the space of theoretical and empirical choices that characterize the mainstream of the shared mental models literature.

Mathieu et al. (2000) considered team members' mental models as comprising knowledge of four separate domains: (1) technology (essentially the mental models described in part one of this chapter); (2) job or task; (3) team interaction (such as roles, communication channels and information flow), and (4) other teammates' knowledge and attitudes. Knowledge of the last three types would rarely be called a mental model outside this literature, and so straight away we can see a broader and more practical orientation than in individually oriented mental models literatures.

For the purposes of operationalization, the authors suggested that these four categories of knowledge may be treated as two: task related and team related. This binary distinction mirrors a distinction that has been made in terms of team behaviors and communications, which have been considered in terms of a task track and a teamwork track (McIntyre and Salas 1995).

Mathieu and colleagues studied dyads using a PC-based flight simulator. One member of each dyad was assigned to the joystick to control aircraft position. The other was assigned to keyboard, speed, weapon systems, and information gathering. Both members could fire weapons. The experimental procedure incorporated a training phase, including the task and basics of teamwork, and then the flying of six missions, divided into three equally difficult blocks of two, each mission lasting around 10 minutes. Performance on a mission was scored in terms of survival, route following, and shooting enemy planes. Team processes were scored by two independent raters viewing videotapes to assign scores, for example, how well the dyad communicated with each other.

Mental models were measured after each pair of missions. At each measurement point, each individual's task or team mental model was elicited by the completion of a relatedness matrix (one for task, one for team), in which the team member rated the degree to which each pair from a set of dimensions was related. For the task model there were eight dimensions, including diving versus climbing; banking or turning; and choosing airspeed. For the team model there

were seven dimensions, including amount of information and roles and team spirit.

Thus, at each measurement point, participants had to assign numbers between -4 (negatively related, a high degree of one requires a low degree of the other) and $+4$ (positively related, a high degree of one requires a high degree of the other) to each pair of dimensions in each domain. For example, they had to rate the relatedness of diving versus climbing to choosing airspeed, and the relatedness of roles to team spirit. For each team at each time for each model-type a convergence index was calculated by computing a correlation co-efficient (QAP correlation) between the two matrices. The co-efficient could vary from -1 (complete disagreement) to $+1$ (completely shared mental models).

The main findings of this investigation were as follows. Contrary to hypothesis, convergence of mental models did not increase over time; rather it was stable across missions 1 to 3. This runs counter to a major and plausible assumption of the shared mental models program, which is that agreement between team members should increase with extent of communication and collaboration.

Nevertheless, convergence of both task and team models predicted the quality of team process and the quality of performance. Further, the relationship between convergence and performance was fully mediated by quality of team process.

The most natural interpretation of these findings is that team process is supported by shared mental models. In turn, good team processes lead to good performance. According to its authors, this study provided the first clear empirical support for the oft-supposed positive relationship between shared mental models and team effectiveness (Mathieu et al. 2000, p. 280).

As well as being paradigmatic in illustrating the key ideas in the shared mental models literature, this study has several aspects that highlight the range of approaches and the controversy in the field.

First, it is worth considering what particular properties of the task and teams may have contributed to the positive relation between shared mental models and team process and performance. Compared with most situations in which coordination and collaboration are of prime interest, including most situations addressed by CSCW researchers, the teams studied by Mathieu et al. were minimal (two members) and the tasks were very short term and relatively circumscribed. Beyond these obvious remarks, I would add that the division of labor in the task was very “close,” and the workers’ performance was extremely interdependent. Of course, interdependence is the signature of collaborative tasks; nevertheless, a situation in which one person controls airspeed and another controls altitude may make this interdependence more immediate than is the norm.

It is also possible that the relatively circumscribed nature of the task and collaboration contributed to the failure of this study to find evidence for the sharing of mental models increasing across the duration of collaboration.

As just mentioned, although the literature contains many proposals that shared mental models will positively influence

process and performance, there has been much less empirical evidence. Another study of particular relevance to HCI is concerned with the workings of software development teams.

Software development is an ideal scenario for the study of team coordination for several reasons. First, much modern software development is quintessentially team based (Crowston and Kammerer 1998; Curtis, Krasner, and Iscoe 1998; Kraut and Streeter 1995), and relies heavily on the complex coordinations of team members. Secondly, this effort is often geographically dispersed, further stressing collaboration and putting an emphasis on communications technologies. Finally, software development takes place in technologically advanced settings with technologically savvy participants, so that it provides something of a test bed for collaboration and communication technologies.

One study of complex geographically distributed software teams has been reported that partially supports the findings of the Mathieu et al. (2000) study and provided complementary evidence for positive effects of shared mental models on team performance. Espinosa et al. (2002) reported a multimethod investigation of software teams in two divisions of a multinational telecommunications company. The most relevant aspect of their study was a survey of 97 engineers engaged in team projects of various sizes ranging from 2 to 7. Team coordination and shared mental models (SMM) were both measured by simple survey items, followed by posthoc correlational analysis to uncover the relation between shared mental models and team process. As in the Mathieu et al. (2000) study, shared mental models were considered in two categories: task and team. A positive relation between team SMM and coordination was discovered, but the effect of task SMM was not significant.

It is worth being clear about the positive relation and how it was computed. Team SMM was computed for each team by assessing the correlations between each team member’s responses to the team SMM survey items. This index was entered as an independent variable in a multiple regression to predict average reported levels of team coordination. It is, of course, hard to infer any causal relation from such correlational analyses, and one might also wonder about the validity of purely questionnaire-based measures of some of the constructs, yet nevertheless the study is highly suggestive that SMM can have a positive influence in group-work situations far removed from pairs of students interactive with a flight simulator. Additionally, Espinosa et al. (2002) reported an interview study in which respondents confirmed their own belief that SMM contributed positively to project communications and outcomes.

Nevertheless, Espinosa et al. (2002) failed to find any relation between task SMM and team process. It seems to me that, in view of the survey methodology, this would have been the more compelling evidence in favor of the SMM construct. It seems less surprising and perhaps less interesting that there should be a correlation between participants’ survey responses concerning how well they communicated on their team, and, for example, their agreement about which teammates had high knowledge about the project.

Levesque, Wilson, and Wholey (2001) reported a different study of software development teams, using ad hoc student-project groupings to study whether sharing of Team SMM increased over time. They only measured Team SMM, using Likert scale items on which participants signaled amount of agreement or disagreement with statements like, “Most of our team’s communication is about technical issues,” “Voicing disagreement on this team is risky,” or “Lines of authority on this team are clear.” Team SMM was measured by computing correlations among team members of these responses after 1, 2, and 3 months of working on a joint project.

Levesque, Wilson, and Wholey (2001) found that, contrary to their hypothesis, team SMM decreased over time. They argue that this is because projects were managed by a division of labor that required much initial collaboration but meant that later activity was more individual.

There are surely many teamwork situations in which role differentiation is critical for success, and this observation suggested that the most straightforward interpretation of shared mental models is overly simple. Indeed, even in teams that continue to meet, communicate, and collaborate, it may be that role differentiation means that task mental models should not so much be “shared” as “distributed” to allow for effective team performance. (Studies of intimate couples have explored a similar process of specialization of memory functions, under the name “transactive memory,” i.e., Wegner 1987, 1995).

When roles are differentiated, it is no longer important that task knowledge is shared, but rather that individuals’ knowledge about who knows what is accurate. Thus, one would expect team SMMs to support communication and collaboration even in teams with highly differentiated roles. This may explain the previously reviewed findings. In the Mathieu et al. study, the team members’ technical roles remained tightly interdependent, so that both task and team models had to be shared for successful performance. In the Espinosa et al. (2002) study, the technical roles may have been differentiated but the level of communication remained high, so that team SMM affected performance but task SMM did not. In the Levesque et al. study, the teams divided their labor to the extent that communication and collaboration ceased to be necessary (apart, perhaps for some final pooling of results). In this case, we would predict that neither task nor team SMMs would affect performance once the division of labor had been accomplished. No data on performance were reported, but team models became less shared over the course of the projects.

Although there has been quite a sudden flurry of interest in shared mental models, this brief review makes clear that much empirical and conceptual work remains to be done. Of particular relevance to this chapter is the question of what exactly is meant by a mental model in this context.

To date throughout the field, mental models have been considered as semantic knowledge, using traditional associative networks as a representation. Thus, mental models have typically been tapped using simple likert scales or direct questions about the relations (i.e. similarity) between

constructs, analyzed with multidimensional techniques such as pathfinder (for a review of measurement techniques in this field, see Mohammed, Kilmoski, and Rentsch 2000). Because interest has focused on the extent to which knowledge and beliefs are common among team members, these approaches have been useful, allowing quantitative measures of similarity and difference. Nevertheless, compared with the literature on individual mental models, they tend to reduce participants’ understanding to mere associations, and yet the thrust of the individual work shows that this may not be appropriate, because the particular conceptualizations of the domain, the analogies drawn, the computational as well as informational relations between internal and external representations, and so on can have real effects on performance. It seems that an important path of development may be to adopt this more refined cognitive orientation and investigate the impact of shared models—as opposed to shared networks of facts and associations—on collaboration.

REFERENCES

- Baddeley, A. 1992. Working memory. *Science* 255:556–9.
- Baddeley, A. D., and D. J. A. Longman. 1978. The influence of length and frequency of training session on the rate of learning to type. *Ergonomics* 21:627–35.
- Bibby, P. A., and S. J. Payne. 1993. Internalization and use—specificity of device knowledge. *Hum Comput Interact* 8:25–56.
- Bibby, P. A., and S. J. Payne. 1996. Instruction and practice in learning about a device. *Cognitive Sci* 20:539–78.
- Bjork, R. A., and T. W. Allen. 1970. The spacing effect: Consolidation or differential encoding? *J Verbal Learn Verbal Behav* 9:567–72.
- Bransford, J. D., J. R. Barclay, and J. J. Franks. 1972. Sentence memory: A constructive versus interpretive approach. *Cogn Psychol* 3:193–209.
- Carroll, J. M. 1990. *The Nurnberg Funnel: Designing Minimalist Instruction for Practical Computer Skill*. Cambridge, MA: MIT Press.
- Carroll, J. M., and J. R. Olson. 1988. Mental models in human–computer interaction. In *Handbook of Human–Computer Interaction*, ed. M. Helander, 45–65. New York: Elsevier.
- Chandler, P., and J. Sweller. 1991. Cognitive load theory and the format of instruction. *Cogn Instruct* 8:293–332.
- Clark, H. H. 1992. *Arenas of Language Use*. Chicago, IL: Chicago University Press.
- Clark, J. M., and A. Paivio. 1991. Dual coding theory and education. *Educ Psychol Rev* 3(3):149–70.
- Cockburn, A., and S. Jones. 1996. Which way now? Analysing and easing inadequacies in WWW navigation. *Int Hum Comput Stud* 45:195–30.
- Collins, A., and D. Gentner. 1987. How people construct mental models. In *Cultural Models in Language and Thought*, ed. D. Holland and N. Quinn. Cambridge, UK: Cambridge University Press.
- Craik, K. J. W. 1943. *The Nature of Explanation*. Cambridge: Cambridge University Press.
- Crowston, K., and E. E. Kammerer. 1998. Coordination and collective mind in software requirements development. *IBM Syst J* 37(2):227–45.
- Curtis, B., H. Krasner, and N. Iscoe. 1988. A field study of the software design process for large systems. *Commun ACM* 31(11):1268–86.

- Diehl, V. A., and C. B. Mills. 1995. The effects of interaction with the device described by procedural text on recall, true/false, and task performance. *Mem Cogn* 23(6):675–88.
- Duggan, G. B., and S. J. Payne. 2001. Interleaving reading and acting while following procedural instructions. *J Exp Psychol Appl* 7(4):297–307.
- Ericsson, K. A., R. T. Krampe, and C. Tesch-Romer. 1993. The role of deliberate practice in the acquisition of expert performance. *Psychol Rev* 100(3):363–406.
- Espinosa, J. A., R. E. Kraut, S. A. Slaughter, J. F. Lerch, J. D. Herbsleb, and A. Mockus. 2002. Shared mental models, familiarity and coordination: A multi-method study of distributed software teams. In *Proceedings of the 23rd International Conference in Information Systems (ICIS)*, 425–33. Barcelona, Spain.
- Fletcher, C. R., and S. T. Chrysler. 1990. Surface forms, textbases and situation models: Recognition memory for three types of textual information. *Discourse Process* 13:175–90.
- Gentner, D., and A. L. Stevens. 1983. *Mental Models*. Hillsdale, NJ: Erlbaum.
- Glenberg, A. M., and W. E. Langston. 1992. Comprehension of illustrated text: Pictures help to build mental models. *J Mem Lang* 31:129–51.
- Halasz, F. G., and T. P. Moran. 1983. Mental models and problem-solving in using a calculator. In *Proceedings of CHI 83 Human Factors in Computing Systems*. New York: ACM.
- Holland, J. H., K. J. Holyoak, R. E. Nisbett, and P. R. Thagard. 1986. *Induction*. Cambridge, MA: MIT Press.
- Johnson-Laird, P. N. 1983. *Mental Models*. Cambridge, UK: Cambridge University Press.
- Johnson-Laird, P. N. 1989. Mental models. In *Foundations of Cognitive Science*, ed. M. I. Posner. Cambridge, MA: MIT Press.
- Kerr, R., and B. Booth. 1978. Specific and varied practice of a motor skill. *Percept Mot Skills* 46:395–401.
- Kieras, D. E., and S. Bovair. 1984. The role of a mental model in learning to use a device. *Cogn Sci* 8:255–73.
- Kintsch, W. 1998. *Comprehension*. Cambridge, UK: Cambridge University Press.
- Kraut, R. E., and L. A. Streeter. 1995. Coordination in software development. *Commun ACM* 38(3):69–81.
- Landauer, T. K., and R. A. Bjork. 1978. Optimum rehearsal patterns and name learning. In *Practical Aspects of Memory*, ed. M. M. Gagneberg, P. E. Morris, and R. N. Sykes, 625–32. London: Academic Press.
- Larkin, J. H., and H. A. Simon. 1987. Why a diagram is (sometimes) worth ten thousand words. *Cogn Sci* 11:65–100.
- Lee, T. D., and R. A. Magill. 1983. The locus of contextual interference in motorskill acquisition. *J Exp Psychol: Learn Mem Cogn* 9:730–46.
- Levesque, L. L., J. M. Wilson, and D. R. Wholey. 2001. Cognitive divergence and shared mental models in software development project teams. *J Organ Behav* 22:135–44.
- Mathieu, J. E., T. S. Heffner, G. F. Goodwin, E. Salas, and J. A. Cannon-Bowers. 2000. The influence of shared mental models team process and performance. *J Appl Psychol* 85:273–83.
- Mayer, R. E. 1996. Learning strategies for making sense out of expository text: The SOI model for guiding three cognitive processes in knowledge construction. *Educ Psychol Rev* 8:357–71.
- Mayer, R. E. 1999a. Research-based principles for the design of instructional messages: The case of multimedia explanations. *Doc Des* 1:7–20.
- Mayer, R. E. 1999b. Multimedia aids to problem solving transfer. *Int J Educ Res* 31:611–23.
- Mayer, R. E. 1999c. Multimedia aids to problem-solving transfer. *Int J Educ Res* 31:661–24.
- Mayer, R. E., and R. B. Anderson. 1991. Animations need narrations: An experimental test of a dual-coding hypothesis. *J Educ Psychol* 83: 484–90.
- Mayer, R. E., and R. B. Anderson. 1992. The instructive animation: Helping students build connections between words and pictures in multimedia learning. *J Educ Psychol* 84:444–52.
- Mayer, R. E., and J. K. Gallini. 1990. When is an illustration worth ten thousand words? *J Educ Psychol* 82:715–26.
- Mayer, R. E., and R. Moreno. 1998. A split-attention affect in multimedia learning: Evidence for dual processing systems in working memory. *J Educ Psychol* 90:312–20.
- Mayer, R. E., and R. Moreno. 2002. Aids to computer-based multimedia learning. *Learn Instruct* 12:107–19.
- Mayer, R. E., and V. K. Sims. 1994. For whom is a picture worth a thousand words? Extensions of a dual-coding theory of multimedia learning. *J Educ Psychol* 86:389–401.
- Mayhew, D. J. 1992. *Principles and Guidelines in Software User Interface Design*. Englewood Cliffs, NJ: Prentice Hall.
- McCloskey, M. 1983. Naïve theories of motion. In *Mental Models*, ed. D. Gentners, and A. L. Stevens, 299–323. Hillsdale, NJ: Erlbaum.
- McIntyre, R. M., and E. Salas. 1995. Measuring and managing for team performance: Emerging principles from complex environments. In *Team Effectiveness and Decision Making in Organizations*, ed. R. A. Guzzo and E. Salas, 9–45. San Francisco, CA: Jossey-Bass.
- McNamara, D. S., E. Kintsch, N. B. Songer, and W. Kintsch. 1996. Are good texts always better? Text coherence, background knowledge, and levels of understanding in learning from text. *Cogn Instruct* 14:1–43.
- Moheammet, S., and B. C. Dunville. 2001. Team mental models in a team knowledge framework: Expanding theory and measurement across disciplinary boundaries. *J Organ Behav* 22(2):89.
- Mohammed, S., R. Klimoski, and J. R. Rentsch. 2000. The measurement of team mental models: We have no shared schema. *Organ Res Methods* 3(2):123–65.
- Moray, N. 1999. Mental models in theory and practice. In *Attention and performance XVII*, ed. D. Gopher and A. Koriati, 223–58. Cambridge, MA: MIT Press.
- Newell, A., and H. A. Simon. 1972. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice Hall.
- Norman, D. A. 1983. Some observations on mental models. In *Mental Models*, ed. D. Gentner and A. L. Stevens, 7–14. Hillsdale, NJ: Erlbaum.
- O'Hara, K. P., and S. J. Payne. 1998. The effects of operator implementation cost on planfulness of problem solving and learning. *Cogn Psychol* 35:34–70.
- O'Hara, K. P., and S. J. Payne. 1999. Planning and user interface: The effects of lockout time and error recovery cost. *Int Hum Comput Stud* 50:41–59.
- Paivio, A. 1986. *Mental Representations*. New York: Oxford University Press.
- Payne, S. J. 1991. A descriptive study of mental models. *Behav Inf Technol* 10:3–21.
- Payne, S. J. 2003. Users' mental models of devices: The very ideas. In *HCI Models, Theories and Frameworks: Towards a Multi-disciplinary Science*, ed. J. M. Carroll, 135–56. San Francisco, CA: Morgan Kaufmann.
- Payne, S. J., A. Howes, and E. Hill. 1992. Conceptual instructions derived from an analysis of device models. *Int J Hum Comput Interact* 4:35–58.

- Payne, S. J., and W. R. Reader. 2006. Constructing structure maps of multiple on-line texts. *Int J Hum Comput Stud* 64(5):461–74.
- Payne, S. J., H. R. Squibb, and A. Howes. 1990. The nature of device models: The yoked state space hypothesis and some experiments with text editors. *Hum Comput Interact* 5:415–44.
- Perfetti, C. A. 1997. Sentences, individual differences, and multiple texts. Three issues in text comprehension. *Discourse Process* 23:337–55.
- Radvansky, G. A., D. H. Spieler, and R. T. Zacks. 1993. Mental model organization. *J Exp Psychol Learn Mem Cogn* 19:95–114.
- Radvansky, G. A., and R. T. Zacks. 1991. Mental models and fact retrieval. *J Exp Psychol Learn Mem Cogn* 17:940–53.
- Rouet, J. F., M. Favart, M. A. Britt, and C. A. Perfetti. 1997. Studying and using multiple documents in history: Effects of discipline expertise. *Cogn Instruct* 75(1):85–106.
- Schmidt, R. A., and R. A. Bjork. 1992. New conceptualizations of practice: Common principles in three paradigms suggest new concepts for training. *Psychol Sci* 3:207–17.
- Schmidt, R. A., D. E. Young, S. Swinnen, and D. C. Shapiro. 1989. Summary knowledge of results for skill acquisition: Support for the guidance hypothesis. *J Exp Psychol Learn Mem Cogn* 15:352–9.
- Shea, J. B., and R. L. 1979. Contextual interference effects on the acquisition, retention, and transfer of a motor skill. *J Exp Psychol Hum Learn Mem* 5:179–87.
- Simon, H. A. 1955. A behavioral model of rational choice. *Q J Econ* 69:99–118.
- Simon, H. A. 1978. On the forms of mental representation. In *Minnesota Studies in the Philosophy of Science*, ed. C. W. Savage, 3–18, Vol. 9. Minneapolis, MN: University of Minnesota Press.
- Simon, H. A. 1992. What is an “explanation” of behavior? *Psychol Sci* 3:150–61.
- Stahl, S. A., C. R. Hind, B. K. Britton, M. M. McNish, and D. Bosquet. 1996. What happens when students read multiple source documents in history. *Read Res* 31(4):430–56.
- Tognazzini, B. 1992. *Tog on Interface*. Reading, MA: Addison-Wesley.
- van Dijk, T. A., and W. Kintsch. 1983. *Strategies of Discourse Comprehension*. New York: Academic Press.
- Wegner, D. M. 1987. Transactive memory: A contemporary analysis of the group mind. In *Theories of Group Behaviour*, ed. I. B. Mullen and G. R. Goethals, 185–208. New York: Springer-Verlag.
- Wegner, D. M. 1995. A computer network model of human transactive memory. *Soc Cogn* 13:319–39.
- Young, R. M. 1983. Surrogates and mappings. Two kinds of conceptual models for interactive devices. In *Mental Models*, ed. D. Gentner and A. L. Stevens, 35–52. Hillsdale, NJ: Erlbaum.

4 Task Loading and Stress in Human–Computer Interaction

Theoretical Frameworks and Mitigation Strategies

James L. Szalma, Gabriella M. Hancock, and Peter A. Hancock

CONTENTS

4.1	Traditional Approaches to Stress Research	56
4.2	Theoretical Frameworks	57
4.2.1	Appraisal Theory	57
4.2.2	Adaptation under Stress	57
4.2.3	The Cognitive–Energetic Framework	59
4.3	Relationship between Stress and Cognitive Workload	60
4.3.1	Cognitive Workload as a Form of Stress	60
4.3.2	Performance and Workload: Associations/Dissociations	60
4.4	Mitigation of Stress	60
4.4.1	Theoretical Bases of Emotion (Stress) Regulation	61
4.4.1.1	Psychological and Physiological Strategies for Emotion Regulation	61
4.4.1.2	Deliberate Emotion Regulation	61
4.4.1.3	Cognitive Reappraisal	62
4.4.1.4	Expressive Suppression	62
4.4.2	Changing the Person	62
4.4.2.1	Training/Skill Development	62
4.4.2.2	Personnel Selection	62
4.4.3	Changing the Task	63
4.4.3.1	Display Design	63
4.4.3.2	Adaptive Automation	63
4.4.4	Hedonomics: Promoting Enjoyable Human–Computer Interaction	63
4.4.4.1	Teletic Work and Intrinsic Motivation	64
4.4.4.2	Intrinsic Motivation and Skill Development	64
4.4.4.3	Teletic Work and Motivational Affordances	64
4.5	Directions for Future Human–Computer Interaction Research in Relation to Workload and Stress	65
4.5.1	Understanding Mental Resources	65
4.5.1.1	Resource Metaphors	65
4.5.1.2	Function of Resources	66
4.5.1.3	Empirical Tests of the Model	66
4.5.1.4	Structural Mechanisms	67
4.5.1.5	Energetic Mechanisms	68
4.5.1.6	Structure and Energy	68
4.5.1.7	Application of Neuroergonomics	68
4.5.2	Development of Adaptation under the Stress Model	68
4.5.2.1	Quantify the Task Dimensions	68
4.5.2.2	Attentional Narrowing	68
4.5.2.3	Stressor Characteristics	69
4.5.2.4	Potential Benefits of Space–Time Distortion	69
4.5.2.5	Changes in Adaptation: The Roles of Time and Intensity	69

4.5.3	Understanding Performance–Workload Associations/Dissociations	69
4.5.3.1	Task Factors	69
4.5.3.2	Multidimensionality of Workload.....	69
4.6	Individual Differences in Performance, Workload, and Stress	70
4.6.1	Trait Differences	70
4.6.2	Affective State Differences.....	70
4.6.2.1	Attentional Narrowing and Adaptive Response	71
4.6.2.2	Hedonomics and Individual Differences	71
4.7	Implications of Stress for Researchers and Practitioners	71
4.8	Summary and Conclusions	72
	References.....	72

Individuals whose professional lives revolve around human–computer interaction (HCI) might well ask themselves why they should even glance at a chapter on stress. It is evident that many computer systems have to support people operating in stressful circumstances and, of course, there are important design issues concerning how to present information in these very demanding circumstances. However, one can legitimately question whether such issues are of any interest to those operating in mainstream HCI. Indeed, if these were the only issues we would most probably agree and recommend the reader to pass on quickly to something of much more evident relevance. However, we hope to persuade the reader that the various aspects of stress research and its application to HCI are not limited to such concerns alone. Indeed, we hope to convince the reader that stress, in its critical form of task loading, is central to all HCIs. To achieve this goal, we first present a perspective that puts stress front and center in the HCI realm. Traditionally, stress has been considered to result from exposure to some adverse environmental circumstances such as excessive heat, cold, noise, or vibration (Hancock, Ross, and Szalma 2007; Conway, Szalma, and Hancock 2007). Its effects manifest themselves primarily in relation with physiological responses most perturbed by the stress at hand. However, Hancock and Warm (1989) observed that stress effects are virtually all mediated through the brain; but for the cortex such effects are almost always of secondary concern since the brain is primarily involved with the goals of ongoing behavior or more simply with dealing with the current task at hand (see Hancock 2010). Therefore, we want to change the orientation of concern here so that stress is not just a result of external interference but rather the primary source of stress comes from the *ongoing task itself*. As we now view the task itself as the primary driving influence, then stress concerns are manifestly and evidently *central to all HCI issues*.

It is one of the clearest paradoxes of modern work that computer-based systems designed to reduce task complexity and cognitive workload actually often impose even greater demands and stresses on the very individuals they are supposed to be helping. Think of how many times in your own work that the computer has appeared to be a barrier to task completion rather than a helpful tool. How individuals cope with such stress has both immediate and prolonged effects on their performance and well-being. Although operational

environments and their associated tasks vary considerably (e.g., air traffic control, baggage screening, hospital patient monitoring, power plant operations, command and control, banking/finance, and general office work), there are certain mechanisms that are common to all stress appraisals and thus to all task demands. Consequently, there are design and HCI principles to address the stress of task demand that can be generalized across many, if not all, domains (Hancock and Szalma 2003a,b). In this chapter we explore these principles to further understand and even exploit stress effects in the HCI domain.

The structure of this chapter flows from these fundamental observations: First, we provide the reader with a brief overview of stress theory and its historical development to set our observations in context. Then we articulate areas for future research, which is needed to more completely understand how stress and workload impact HCI and how their positive effects can be exploited while mitigating their negative effects. We conclude by providing an overview of these principles and some directions for future effort.

4.1 TRADITIONAL APPROACHES TO STRESS RESEARCH

As we have seen, stress has traditionally been conceived of as either an external, aversive stimulus (constituted of physical, cognitive, or social stimulation patterns) imposed on an individual or that person's individual response to such perturbations. Each of these theoretical perspectives has limited explanatory power. Considering stress as an external stimulation is useful for categorizing effects of physical environments (e.g., heat, noise, vibration), but such an approach cannot explain why the same stimulus pattern produces vastly different effects on different individuals. Physiological interpretations (e.g., Selye 1976) have utilized arousal explanations of stress. However, more recent demonstrations that different sources of stress are associated with different patterns of cognitive effects make it clear that adaptation or so-called arousal theories of stress cannot, by themselves, completely address the issue (Hockey 1984; Hockey and Hamilton 1983; Hockey, Gaillard, and Coles 1986).

Thus, to understand stress effects we now have to embrace an even wider, multidimensional perspective (e.g., Matthews

2001). In this chapter, we choose to emphasize a view of stress as primarily an outcome of the appraisal of environmental demands that either tax or exceed an individual's resources to cope with that demand. These person–environment *transactions* (Lazarus and Folkman 1984) occur at multiple levels within an organism (Matthews 2001; van Reekum and Scherer 1997). Further, these processes represent efforts by the organism to adapt to imposed demands via regulation of its own internal state while seeking to change the external environment (e.g., obtaining shelter). In Section 4.2, we describe the theoretical frameworks that guide our observations on stress in the context of HCI. These perspectives emerge from the work of Lazarus (1999; and see also Lazarus and Folkman 1984), Hancock and Warm (1989), and Hockey (1997).

4.2 THEORETICAL FRAMEWORKS

Herein is a brief introduction to key theories regarding stress and its effects on performance.

4.2.1 APPRAISAL THEORY

Among the spectrum of cognitive theories of stress and emotion, perhaps the best known is the “cognitive–motivational–relational” theory proposed by Richard Lazarus and his colleagues (see Lazarus 1991, 1999; Lazarus and Folkman 1984). This theory is cognitive in that stress and emotion each depends on an individual's cognitive appraisals of internal and external events. These appraisals in their turn depend in part on the person's knowledge and experience (cf. Bless 2001). The theory is motivational in that emotions in general, including stress responses, are reactions to one's perceived state of progress toward or away from one's goals (see Carver and Scheier 1998). The relational aspect emphasizes the importance of the transaction between individuals and their environment. Together these three components shape the emotional and stress state of an individual. The outcomes of these processes are patterns of appraisal that Lazarus (1991) refers to as “core relational themes.” For instance, the core relational theme for anxiety is uncertainty and existential threat, whereas that for happiness is evident progress toward goal achievement. Thus, when individuals appraise events relative to their desired outcomes (goals), negative, “goal-incongruent” emotions and stress can be produced if such events are appraised as hindering progress. Conversely, promotion of well-being and pleasure occurs when events are appraised as facilitating progress toward a goal (i.e., goal-congruent emotions). Promotion of pleasure and happiness (see Hancock, Pepe, and Murphy 2005; Ryan and Deci 2001), therefore, requires the design of environments and tasks themselves that afford goal-congruent emotions. The understanding of interface characteristics in HCI that facilitate positive appraisals and reduce negative appraisals is thus a crucial issue and an obvious avenue in which HCI and stress research can fruitfully interact.

A major limitation of all appraisal theories, however, is neglecting to understand how task parameters influence resulting coping response. Although the appraisal mechanism itself may be similar across individuals and contexts (e.g., see Scherer 1999), the specific content (e.g., which events are appraised as a threat to well-being) obviously varies across individuals and contexts. One would expect that the criteria for appraisal (e.g., personal relevance, self-efficacy for coping) are similar across individuals for specific task parameters as for any other stimulus or event. However, individual differences occur in the specific content of an appraisal (e.g., one person's threat is another's challenge) and therefore in the resultant response. An understanding of stress effects in HCI thus requires understanding the task and person factors and treating the transaction between the human being and the system as the primary unit of analysis (see Lazarus and Folkman 1984). This entails knowing how different individuals appraise specific task parameters and how changes in knowledge structures might ameliorate negative stress effects and promote adaptive affective states in human–technology interaction. A visual representation of this emergent unit of analysis that comes from the interaction of a person and the environment, including the task, is shown in Figure 4.1 (Hancock 1997).

4.2.2 ADAPTATION UNDER STRESS

A theoretical framework developed specifically for stress as it relates to performance is the maximal adaptability model presented by Hancock and Warm (1989). They distinguished three facets of stress and labeled them the “trinity of stress,” as shown in Figure 4.2. Input refers to the environmental events to which an individual is exposed, which include information (i.e., displays) as well as traditional input categories such as temperature, noise, and vibration (e.g., Hancock, Ross, and Szalma 2007; Pilcher et al. 2002). The second is adaptation, which encompasses the appraisal mechanisms referred to previously. The third and final component is output level, which indicates how an organism behaves with respect to goal achievement. A fundamental tenet of the Hancock and




FIGURE 4.1 An illustration of the emergence of a supraordinate unit of analysis that derives from the interaction of an individual (person), the tool he or she uses (computer), the task he or she has to perform (task), and the context (environment) against which the action occurs.

Warm (1989) model is that in a large majority of situations (and even in situations of quite high demand) individuals do adapt effectively to input disturbance. That is, they can tolerate high levels of either overload or underload without enormous change to their performance capacity. Adaptive processes occur at multiple levels, some being the physiological, behavioral (i.e., performance), and subjective/affective levels. These adaptations are represented in the model as a series of nested, extended inverted-U functions (see Figure 4.3) that reflect the fact that under most conditions the adaptive state of the organism is stable. However, under extremes of environmental underload or overload, “failures” in adaptation do occur. Thus, as the individual is perturbed by the input, the first threshold one traverses is subjective




FIGURE 4.2 The trinity of stress, which identifies three possible “loci” of stress. It can be viewed as an input from the physical environment, which can be described deterministically. Since such a profile is by definition unique, it is referred to as a stress signature. The second locus is adaptation, which describes the populational or nomothetic reaction to the input itself. It is most evidently measurable in the processes of compensation. The third and final locus is output, which is expressed as the impact on the ongoing stream of behavior. Since the goals of different individuals almost always vary, this output is largely idiographic or person specific. It is this facet of stress that has been very much neglected in prior and contemporary research.

comfort. This is followed by behavioral effects and finally failure of the physiological system (e.g., loss of consciousness). Examples of such extreme failures are relatively rare in most settings, although when they do occur (e.g., in conflict situations) they are often catastrophic for the individual and the system he or she is operating (e.g., Harris, Hancock, and Harris 2005).

This model is unique in that it provides explicit recognition that the *proximal form of stress* in almost all circumstances is the task itself. Task characteristics are incorporated in the model by two distinct base axes representing spatial and temporal components of any specified task. Information structure (the spatial dimension) represents how task elements are organized, including challenges to such psychological capacities such as working memory, attention, decision making, and response capacity. The temporal dimension is represented by information rate. Together, these dimensions can be used to form a vector (see Figure 4.4) that serves to identify the current state of adaptation of the individual. Thus, if the combination of task characteristics and an individual’s stress level can be specified, a vector representation can be used to predict behavioral and physiological adaptation. The challenge lies in quantifying the information-processing components of cognitive work (see Hancock, Szalma, and Oron-Gilad 2005).

Although the model shown in Figure 4.4 describes the level of adaptive function, it does not articulate the mechanisms by which such adaptation occurs. Hancock and Warm (1989) argued that one way in which individuals adapt to stress is by narrowing their attention by excluding task-irrelevant cues (Easterbrook 1959). Such effects are known to occur in




FIGURE 4.3 The extended-U relationship between stress level and response capacity. As is evident, the form of degradation is common across the different reflections of response. At the center of the continuum is the normative zone, which reflects optimal functioning. Outside of this is the comfort zone, which reflects the behavioral recognition of a state of satisfaction. Beyond this lies the reaction of psychological or cognitive performance capacity. Finally, the outer envelope is composed of physiological functioning. There are proposed strong linkages between the deviation from stability at one level being matched to the onset of radical failure at the more vulnerable level that is nested within it. The model is symmetrical in that underload (hypostress) has mirror effects to overload (hyperstress), which is usually considered the commonly perceived interpretation of stress.




FIGURE 4.4 A three-dimensional representation of Figure 4.3. The description given in Figure 4.3 is now expanded into a three-dimensional representation by parsing the base hypostress–hyperstress axis into its two component elements. These divisions are composed of information rate (the temporal axis) and information structure (the spatial axis). Note that any one source of input stress can be described as a scalar on the base axis and these scalars can be summed to provide a multi-input stress vector that then provides a prediction of both performance and physiological adaptability, which are the primary descriptors on the vertical axis. A, B, C and D represent the thresholds of adaptability. A represents the physiological zone of maximal adaptability; B is the psychological zone of adaptability, C is the comfort zone, while D illustrates the normative zone.

spatial perception (e.g., Bursill 1958; Cornsweet 1969), and narrowing can occur at the levels of both central and peripheral neural systems (Dirkin and Hancock 1984, 1985; Hancock and Dirkin 1983). Further, Hancock and Weaver (2005) have argued that distortions of temporal perception under stress are related to this narrowing effect. However, evidence suggests that these two perceptual dimensions (space and time) may not share common perceptual mechanisms (see Ross et al. 2003; Thropp, Szalma, and Hancock 2004).

4.2.3 THE COGNITIVE–ENERGETIC FRAMEWORK

The Hancock and Warm model accounts for the levels of adaptation and adaptation changes that occur under the driving forces of stress. However, it does not articulate precisely how effort is allocated under stress or the mechanisms by which individuals appraise the task parameters that are the proximal source of stress. The precise effort allocation issue is addressed by a cognitive–energetic framework described by Hockey (1997). The compensatory control model is based on three assumptions: (1) Behavior is goal directed. (2) Self-regulatory processes control goal states. (3) Regulatory activity has energetic costs (i.e., it consumes




FIGURE 4.5 The two-level effort regulation model by Hockey: This model provides a mechanism by which an individual allocates limited cognitive resources to different aspects of performance. (From Hockey, G. R. J., *Biol Psychol*, 45:73–93, 1997. With permission.)

resources). In this model, a compensatory control mechanism allocates resources dynamically according to the goals of an individual and the environmental constraints. The mechanisms operate at two levels (see Figure 4.5): The lower level is more or less “automatic” and represents established skills. Regulation at this level requires few energetic resources or

low active regulation and effort (cf. Schneider and Shiffrin 1977). The upper level is a supervisory controller, which can shift resources (effort) strategically to maintain adaptation, and reflects effortful and controlled processing. The operation of the automatic lower loop is regulated by an effort monitor, which detects changes in the regulatory demands placed on the lower loop. When demand increases beyond the capacity of the lower loop, control is shifted to the higher, controlled processing loop. Two strategic responses of the supervisory system are increase in effort and change in the goals. Goals can be modified in terms of their kind (change the goal itself) or strength (e.g., lowering the criterion for performance). From a self-regulation perspective, these modifications adjust the discrepancy between goal state and the current state by increasing effort or changing the goal (see Carver and Scheier 1998).

4.3 RELATIONSHIP BETWEEN STRESS AND COGNITIVE WORKLOAD

4.3.1 COGNITIVE WORKLOAD AS A FORM OF STRESS

The Hancock and Warm (1989) model explicitly identifies the task itself as the proximal source of stress. In operational environments, this is often manifested as increases or decreases in cognitive workload (Moray 1979). As in the case of stress, workload is easily identified colloquially but difficult to define operationally. Workload can manifest in terms of the amount of information to be processed (an aspect of information structure) and the time available for processing (information rate). Thus, the base axes of the Hancock and Warm model capture dimensions of workload as well as stress (and see Hancock and Caird 1993). Indeed, physiological measures of workload (O'Donnell and Eggemeier 1986) are often the same as the measures used to assess physiological stress. Similarly, subjective measures of workload and stress both reflect appraisals of the task environment and of its perceived effect on the individual (Hart and Staveland 1988). Although the two concepts were developed in separate research traditions, the artificial boundary between them should be dissolved as each term refers to similar processes. The implication for HCI is that computer-based tasks that impose either too much or too little demand will likely be appraised as stressful. In the latter case, the underload stress will be interpreted as boredom. Thus, the design process for the development of computer interfaces should include assessment of perceived workload as well as affective state.

4.3.2 PERFORMANCE AND WORKLOAD: ASSOCIATIONS/DISSOCIATIONS

It is often the case that performance is maintained under increased workload/stress, which is reflected in the extended-U model described by Hancock and Warm and in the mechanisms of Hockey's energetic model of compensatory control. Maintaining performance under stress is associated with costs, both physiologically and cognitively.

Further, one would expect that in easier tasks performance is not as costly and that, therefore, there is a direct *association* between task difficulty and perceived workload. Such performance–workload associations do occur, and they occur most prevalently in vigilance tasks (Warm, Dember, and Hancock 1996; see also Szalma et al. 2004). However, other forms of workload–performance relations can occur. For instance, perceived workload may change as a function of change in task demand, but performance remains constant. Hancock (1996) refers to these situations as *insensitivities*, which can be diagnostic with respect to the relation between the individual and the task (see also Parasuraman and Hancock 2001). Thus, consistent with the frameworks of Hancock and Warm (1989) and Hockey (1997), one response to increased task demand is to exert more effort, thereby maintaining performance but increasing perceived workload. Alternatively, one could have a situation in which task demands increase and performance decreases, but perceived workload does not change. This suggests that appraisals of a task are not always sensitive to actual changes in that task.

Interesting corollaries of these observations are performance–workload *dissociations* that sometimes occur (Hancock 1996; Yeh and Wickens 1988). In such cases, decreased performance is accompanied by *decreased* workload. One possible reason for such a result is disengagement of the individual from the task (i.e., the person gives up; see Hancock 1996). In the case where increased performance is observed to be accompanied by increased perceived workload, the pattern suggests effective improvement of performance at the cost of increased effort allocation. An area of much-needed research is establishing which task parameters control the patterns of performance–workload associations and dissociations, and how these change dynamically as a function of time on task. It may well be that reformulating the task by innovations in the interface itself addresses these crucial concerns (see Hancock 1997). Indeed, the structure and organization of computer interfaces will be a major factor in determining both performance under stress and the relation of performance to perceived workload.

4.4 MITIGATION OF STRESS

If changing the fundamental nature of demand is one solution, we now look at other approaches to mitigate the negative effects of stress and workload on the performance of HCI tasks. These strategies include skill development (e.g., Hancock 1986), specific display design changes (Hancock and Szalma 2003a; Wickens 1996), as well as technologies employing adaptive automation and decision aids (Hancock and Chignell 1987). Developing skills so that they are relatively automatic as opposed to the alternative controlled processing (Schneider and Shiffrin 1977) and developing expertise can mitigate some of the negative effects of stress (Hancock and Hancock 2010). Regarding display design, simple, easily perceivable graphics can permit quick, direct extraction of information when cognitive resources are reduced by stress and workload (Hancock and Szalma

2003a). Adaptive automation can be employed by adjusting the level of automation and the management of that automation according to stress state (e.g., Scerbo, Freeman, and Mikulka 2003). In addition, adapting the form of automation (i.e., level, management type) to the operator based on their own personal style of interaction can serve to improve its utility for aiding performance and reducing stress and workload (see Thropp et al. 2004). Indeed, experimental findings are even now beginning to establish the relation for automation to effectively mitigate performance-related stress (Funke et al. 2007).

4.4.1 THEORETICAL BASES OF EMOTION (STRESS) REGULATION

In both theory and research, stress is clearly linked with the more general topic of emotion (Lazarus 1999). Indeed there is growing recognition of the need to consider a user's emotional response to a task or an interface as it is an important aspect of design. Emotions are valenced reactions to either internal or external stimuli, which trigger multisystemic changes in both physiology and behavior (Ochsner and Gross 2005). Emotions are therefore useful in presenting feedback concerning an operator's ongoing interaction with the environment and especially the computer-based technology with which they must interact (Folkman et al. 1986). The computer and the manner in which it functions can produce a range of emotional reactions in the operator who is attempting to manipulate the system. However, at present the computer has no concept of emotional experience or display, no matter how often humans might attribute these characteristics to the machine (Luczak, Roetting, and Schmidt 2003). Contrary to intuition, the computer is not malfunctioning in order to frustrate or spite its users and no amount of shouting or banging will presently instill it with a sense of motivation to work. However, this is not to say extensive efforts are not underway to develop computer systems that both recognize user emotions and generate emotional expressions on behalf of the computer system (Zhang et al. 2010). The very act of interacting with a computer represents an emotional experience as the machine is a tool by which the operator hopes to accomplish his or her desired goals. Events whereby the computer facilitates the achievement of a goal can result in pleasant emotions (i.e., accomplishment, relief, and happiness), whereas instances in which the interface component of the human–computer dyad is perceived as detrimental or as a barrier to goal fulfillment can produce negative valence emotions such as anger and frustration (Hassenzahl and Ullrich 2007). Therefore, emotions are an inherent feature of HCI as no human action, even one performed in tandem with an affectless instrument, takes place in a completely emotion-free context. Emotions then have the potential in HCI to become either stressors themselves or tools by which operators can cope with stress and enhance the effectiveness and efficiency of performance. Gross, Richards, and John (2006) have postulated that effective emotion regulation is a

qualification for adaptive functioning in almost all everyday skills. Given the growing popularity and availability of technologies such as personal digital assistants and cellular telephones, HCIs are rapidly becoming modal everyday tasks. Techniques for regulating the pervasive influence of emotions are therefore useful skills for the operator to develop so that they may minimize the disadvantageous consequences of negative emotional experiences.

4.4.1.1 Psychological and Physiological Strategies for Emotion Regulation

Emotions are typically categorized based on two componential characteristics: (1) valence and (2) arousal (Lang 1995). Valence refers to the extent to which an operator interprets an emotion as pleasant or unpleasant. Arousal constitutes the extent to which an emotion evokes a response from the operator's physiological system. Stress is unique in that it has the ability to run the gamut on both dimensions; thus, it can be perceived as both pleasant and unpleasant, as well as inducing either mild or severe physiological reactions. Techniques for its regulation therefore incorporate methods to address both the psychological and physiological components of stress (see Hancock and Warm 1989).

A typical course of emotional experience, without any attempts at emotion regulation, begins with emotionally charged environmental cues eliciting intrapersonal emotional response tendencies, which lead to emotional responses (Gross 1998). Emotion regulation techniques may therefore intercede at a number of points in this process. Antecedent-focused strategies, such as cognitive reappraisal, specifically endeavor to manipulate the interaction between emotional cues and their subsequent response tendencies and are therefore a method of evaluation (i.e., viewing the interaction with the computer as a learning experience). Response-focused strategies such as expressive suppression, on the other hand, affect the relationship between an operator's response tendencies and the resultant emotional response and are therefore a method of modulation (i.e., ignoring any unpleasant feelings resulting from interaction with the computer). Both techniques have proved to be effective strategies, although the optimal technique heavily depends on the situation (Gross 2002).

4.4.1.2 Deliberate Emotion Regulation

Emotion regulation entails “processes that individuals use to influence which emotions they generate, when they do so, and how these emotions are experienced or expressed” (Ochsner and Gross 2005, p. 243). Active processing of this nature, which requires attentional resources, is referred to as deliberate emotion regulation. Recent research suggests the possibility that emotion regulation can take place automatically, at an unconscious level (Mauss, Bunge, and Gross 2007). While seminal research efforts investigating the influence of automatic emotion regulation on performance are currently underway (Hancock and Beatty 2010), our focus here necessarily concentrates on the more established deliberate

emotion regulation strategies. The two strategies that are studied most often are cognitive reappraisal and expressive suppression. These techniques share a common goal of emotion regulation, but they differ in the aspects of emotion that they influence, when they begin to influence emotional experience, and their long-term versus short-term effectiveness.

4.4.1.3 Cognitive Reappraisal

Cognitive reappraisal is defined as the act of interpreting potential emotion-provoking stimuli in unemotional terms (Speisman et al. 1964). The purpose of cognitive reappraisal is therefore to influence an individual's cognition, in order to maintain control over emotional responses. As mentioned in Section 4.4.1.1, cognitive reappraisal is an antecedent-focused emotional regulation strategy. The intervention occurs as early as possible in the emotional experience so as to minimize deleterious performance consequences. Indeed, Gross (2002) reported that reappraisal is more effective than expressive suppression. Unlike suppression, it has no detrimental effects on other cognitive processes such as memory. The utilization of cognitive reappraisal is also associated with superior long-term health outcomes (Haga, Kraft, and Corby 2009).

4.4.1.4 Expressive Suppression

Expressive suppression is defined as the inhibition of emotionally expressive behavior despite emotional arousal (Gross and Levenson 1993). The primary aim of the suppression approach is therefore to minimize outward displays of emotion, that is, targeting overt behavior instead of cognition. As a response-focused strategy, suppression initiates its influence later in the process sequence than cognitive reappraisal (Gross 1998). Although cognitive reappraisal is more effective over the long term, suppression is the superior short-term option; operators employing expressive suppression may have more attentional resources available as they are not actively engaging in the continual assessment and reassessment of environmental stimuli. Operators engaging in HCIs under time constraints may therefore find this strategy more effective.

4.4.2 CHANGING THE PERSON

To improve a stressful human-machine interaction, one viable option is to alter the human's attitudes or abilities either through training or selection.

4.4.2.1 Training/Skill Development

Clearly, the greater the skill of an individual the more resilient his or her performance under stress (Hancock 1986). This well-established phenomenon is incorporated into the energetic theories of stress and performance and is an approach most often taken to mitigate adverse workload and stress effects. However, training on relevant tasks is only one method of training for stress. There are also techniques for training individuals to cope more effectively with stress itself, such as the aforementioned emotion regulation techniques

discussed in Sections 4.4.1.3 and 4.4.1.4, which essentially build stress-coping skills. An additional example of such an approach is stress exposure training (SET; Johnston and Cannon-Bowers 1996), a three-phase procedure in which individuals are provided information regarding the stress associated with task performance, are provided training on the task, and then practice their task skills under simulated stress conditions. This technique has been shown to be effective in reducing anxiety and enhancing performance (Saunders et al. 1996). There is evidence that coping skills learned with a particular type of stressor and task can be transferred to novel stressors and tasks (Driskell, Johnston, and Salas 2001). For such an intervention to succeed, however, it is crucial that the training is designed based on a complete and accurate analysis of the task environment (Johnston and Cannon-Bowers 1996). If task parameters that are the most responsible for the workload and stress are identified, these can be especially targeted in training.

An additional issue in training for more effective stress coping is modifying an individual's appraisal of events, an approach that is coincident with the emotion regulation technique of cognitive reappraisal discussed in Section 4.4.1.3. Learning to approach HCIs with effective coping is therefore a valuable skill to acquire as early as possible in any task training. Inducing automaticity in some skills permits reallocation of resources to coping efforts, as well as reducing the likelihood that the task environment itself is appraised as threatening. Even if an event is appraised as a threat to an individual's psychological or physical well-being, the highly skilled individual will appraise his or her coping ability as sufficient to handle such an increased demand. However, there has been limited research on how individuals who develop expertise also develop the capacity to effectively cope with the stress that accompanies performance in a given domain and on the extent to which stress-coping skills in one domain transfer to other domains. Deliberate practice generally facilitates skill development (Ericsson 2006). If one considers coping with stress to be a skill, then in principle deliberate practice should permit the development of expertise in coping with stress. This likely involves parsing the task into components, based on cognitive task analysis, and designing training procedures that target the stressful aspects of the task. However, such efforts require an understanding of how different forms of stress affect different forms of information processing. Since these variables are difficult to quantify, establishing these linkages must be driven by theory. Elucidation of these issues will provide the groundwork for future development of stress mitigation tools during training and skill development.

4.4.2.2 Personnel Selection

Selection techniques have been a popular choice for matching individuals to specific jobs, but the focus has traditionally and historically been on intellectual skills (e.g., Yerkes 1918). Selecting individuals for their stress-coping capability has been applied to the selection criteria for police officers, who therefore tend to be as stable as or more emotionally stable

than the rest of the population (for a review, see the work by Brown and Campbell [1994]). Selecting individuals with proficient stress-coping skills becomes still more difficult given the complex criteria that define stress and the fact that “successful” coping skills vary by situation and desired outcome. Research is therefore needed that links particular traits to stress-coping skills for specific task environments. The effectiveness of everyday life stress coping, such as that observed in individuals who are extraverted (McCrae and Costa 1986; Penley and Tomaka 2002) or optimistic (Aspinwall, Richter, and Hoffman 2002; Scheier and Carver 1985), may not predict effective coping in specific task domains. Understanding which individuals will likely cope effectively with a particular task therefore requires first a thorough understanding of the perceptual, cognitive, and psychomotor characteristics of the task and then linking these parameters to trait profiles. By far the most research on the relation of affective traits to task performance has been conducted in the areas of extraversion and trait anxiety/neuroticism (see the work by Matthews, Deary, and Whiteman [2003] for a review). However, the characteristics of greatest interest may vary somewhat across domains, although some general traits (e.g., emotional stability, conscientiousness) would be expected to moderate performance across a variety of task environments.

4.4.3 CHANGING THE TASK

Modifying the technology component is another possibility for more effective HCI.

4.4.3.1 Display Design

Although training and selection can mitigate stress effects, the primary method of stress mitigation requires the tasks themselves to be redesigned. This is for two reasons: (1) There will be many instances where selection is not possible and expenditure of significant resources on training is undesirable. (2) There are instances in which one wishes to design an interface that requires little or no training and that can be used by any member of a large population of individuals (e.g., consumers). Particularly in light of the observation that task represents the proximal source of stress, future work in stress mitigation for HCI should focus on redesign of the task and of the interface itself. In previous work, we have argued that existing display design techniques that are simple and easily perceived would be the best choice for an interface that is used in stressful environments (Hancock and Szalma 2003a). Specifically, configural or object displays can represent complex, multivariable systems as simple geometric shapes or emergent features if those features are mapped well to system dynamics (see Bennett and Flach 1992). Under stress, it is the complex problem solving and analytical skills that are the most vulnerable and are apt to decline first. A display that allows fast extraction of information with minimal cost in working memory load can mitigate stress effects (Hancock and Szalma 2003a; Wickens 1996). A combination of training to automaticity and displays of information that can be perceived directly with a minimum

of information-processing requirements is currently one of the best approaches for stress mitigation in cognitively complex environments.

4.4.3.2 Adaptive Automation

Another approach for stress mitigation is the allocation of function to automated systems (Hancock and Chignell 1987). The advent of modern automated systems allows for automation to adapt to the state of an individual (Scerbo, Freeman, and Mikulka 2003). Thus, at points in time when an operator is overstressed and overtaxed, the system can assume control of some task functions, thereby freeing resources to effectively cope with increased task demand. Two potential problems for automated systems are that overreliance can occur and operator skills can atrophy. However, a dynamic (adaptive) automated system that permits or requires the operator to perform functions at different points in time can reduce the probability of skill atrophy while still relieving the workload and stress of task performance.

However, the introduction of automation can itself induce stress. Operators who work with automated systems, particularly static, inflexible automated systems, are relegated to the role of monitors who must respond only when untoward events occur. Sustained attention requirements are in fact quite stressful (Warm, Parasuraman, and Matthews 2008) and paradoxically induce high perceived workload (Warm, Dember, and Hancock 1996). Adaptive automation can mitigate this problem by dynamically assigning tasks to the machine or the human being depending on the environmental conditions and the state of the operator (Hancock and Chignell 1987). Indeed, potential techniques for enabling the operator’s neurological state to adjust to automation have been identified (e.g., Scerbo, Freeman, and Mikulka 2003).

4.4.4 HEDONOMICS: PROMOTING ENJOYABLE HUMAN–COMPUTER INTERACTION

Stress research has traditionally followed the edict of ergonomics and human factors, in general, to first do no harm and then seek to prevent pain and injury. As with the rest of behavioral science, stress researchers have often sought to treat the symptoms of stress and mitigate its negative effects on performance. However, with the advent of positive psychology (Seligman and Csikszentmihalyi 2000), there has been a movement to incorporate the promotion of pleasure and well-being rather than restrict efforts to pain prevention. Hancock (Hancock, Pepe, and Murphy 2005) coined the term hedonomics and defined it as that branch of science that facilitates the pleasant or enjoyable aspects of human–technology interaction. In short, the goal of hedonomics is to design with happiness in mind. Hedonomics is a fairly new research area, but during the last decade there has been a rapid growth in research concerning affect and pleasure. Affective evaluations provide a new and different perspective in human factors engineering. It is not how to evaluate users; it is how the user evaluates (Hancock et al. 2005). The research on hedonic values and seductive interfaces is in fact a welcome contrast

to research on safety and productivity, which have dominated human factors and ergonomics (HF/E) for so long. It must be noted, however, that *pleasurable* interaction with technology is not necessarily conducive to *happiness*. Indulging in pleasures can sometimes interfere with happiness and well-being (see Fromm 1976; Kasser 2002; Ryan and Deci 2001).

Our argument is not that we should discard current methods in HF/E. Clearly functionality and usability are necessary conditions for pleasurable interaction with technology. If an interface does not function in a way that is congruent with the user's goals so that the user appraises the technology as an agent that is interfering with goal achievement, that interaction is likely to be stressful and performance may decline. However, functionality and usability are necessary but not sufficient conditions for pleasurable interactions with technology. The interface should be designed such that it affords appraisals of the technology as a convivial tool (Illich 1973) or aid. One can also utilize the human tendency to anthropomorphize technology to facilitate such appraisals of the technology as “helpful and supportive” rather than as “conflictive” or, worse, an “enemy” (Luczak, Roetting, and Schmidt 2003).

Hedonomic design is of obvious importance for the development of consumer products, but in principle it can also transform the very nature of work itself, rendering it “fun.” Although there may be some tasks that will never be completely enjoyable, there are many individuals who have jobs that could be made more enjoyable by designing the tasks such that they promote teleetic work (Csikszentmihalyi 1990) while also facilitating intrinsic motivation (Deci and Ryan 2000).

4.4.4.1 Teletic Work and Intrinsic Motivation

A useful theoretical framework for hedonomics is self-determination theory (SDT; Deci and Ryan 1985, 2000; Ryan and Deci 2000, 2001). From this perspective, there are three organismic needs that are essential for facilitating intrinsic motivation for task activity and the positive affect that can accompany such states. These needs are competence (self-efficacy; see also Bandura 1997), autonomy (personal agency, not independence per se), and relatedness. An important difference between this theory and other theories of motivation is the recognition that there are qualitatively different forms of motivation (Gagne and Deci 2005). Thus, in SDT five categories of motivated behavior are identified that vary in the degree to which motivation is self-determined. Four of the categories reflect extrinsic motivation and one category is intrinsic motivation. In the latter case, individuals are inherently motivated to engage in activity for its own sake or for novelty and challenge. The four extrinsic motivation categories vary in the degree to which regulation of behavior is internalized by the individual and, therefore, they are more autonomous and self-determined (Ryan and Deci 2000). The process of internalization involves transforming an external regulation or value into one that matches an individual's own values. The development of such autonomous motivation is crucial to skill development, since the person must maintain his or her effort throughout a long and arduous process. Individuals who are autonomously motivated to learn are those who

develop a variety of effective self-regulation strategies, have high self-efficacy, and set a number of goals for themselves (Zimmerman 2000). Further, effective self-regulation develops in four stages: (1) observation, (2) emulation, (3) self-control, and (4) self-regulation. Successful skill development involves focus on process goals in the early stages of learning and outcome goals in the fourth stage (Zimmerman 2000).

4.4.4.2 Intrinsic Motivation and Skill Development

Research has established that intrinsic motivation is facilitated by conditions promoting autonomy, competence, and relatedness (see Deci and Ryan 2000). Three factors that support autonomy are as follows: (1) meaningful rationales for doing a task, (2) acknowledgment that the task might not be interesting, and (3) an emphasis on choice rather than control. It is important to note that externally regulated motivation predicts poorer performance on heuristic tasks (Gagne and Deci 2005), suggesting that as experts develop better knowledge representations it will be crucial to promote their internal regulation of motivation. Although intrinsic motivation has been linked to how task activities and environmental contexts meet psychological needs, it is not clear why skilled performers are able to meet these needs or why an individual chooses a particular computer interface. It is likely that interest in activities codevelops with abilities and traits (see Ackerman and Heggstad 1997), but this issue needs more thorough investigation in the context of complex computer environments that require highly skilled work.

Emotions can be powerful motivators of task performance, including tasks involving HCI. Both intrinsic motivation and emotional experience play critical roles in beginners' perceptions concerning their current and future interactions with a system (Venkatesh 2000). Although emotion cannot and most probably should not be designed out of an HCI, it is possible to design activities in which emotional experience facilitates learning and enhances an operator's intrinsic motivation for task mastery (Lepper and Cordova 1992). Such activities should also aim to simultaneously foster effective emotion regulation techniques. Learning to perform in the presence of common stressors early in the learning process will help to maintain task engagement. In addition to the aforementioned concern and the issues of efficacy and self-regulation, there is a need to examine the process by which individuals internalize extrinsic motivation as gain experience with a particular interface or system. In particular, Gagne and Deci (2005) noted that little research has examined the effect of reward structures and work environments on the internalization process. It is likely that environments structured to meet basic needs more likely facilitate internalization processes and inoculate learners against the trials and tribulations that face them as they interact with new technologies.

4.4.4.3 Teletic Work and Motivational Affordances

Teletic, or autoteletic, work refers to work that is experienced as enjoyable and is associated with “flow” or optimal experience characterized by a sense of well-being and harmony with one's surroundings (Csikszentmihalyi 1990). There is variation

in both tasks and individuals with respect to the degree to which the human–technology interaction is teleetic. There are four categories in which individuals tend to fall with respect to their relation to work: First, there is a small proportion of the population that is always happy in life, regardless of their activity. Csikszentmihalyi (1990) refers to these individuals as having an “autotelic personality.” There is also a group of individuals who are naturally predisposed to being happy regarding a specific task. They appraise such tasks as enjoyable and often seek out these activities. The third group consists of individuals who enjoy specific activities but cannot do them professionally. This group includes individuals such as amateur athletes. The vast majority of people, however, do work for purely functional reasons (e.g., finances and security). For these individuals, work is boring and grinding because the task itself is nearly always considered aversive. A goal of hedonomics is to design work that can be enjoyed to the greatest extent possible. This means structuring the environment as an entire system, ranging from the specific cognitive and psychomotor demands to the organization in which a person works. Even in jobs that are not inherently enjoyable, some degree of positive affect can be experienced by workers if their environment is structured to facilitate a sense of autonomy (personal agency), competence, and relatedness (Deci and Ryan 2000; see also Gagne and Deci 2005). From an ecological perspective (Flach et al. 1995), this means identifying the *motivational affordances* in the task and work environment, and designing for these affordances. Thus, just as one might analyze the affordance structure of an interface using ecological interface design methods (e.g., Vicente and Rasmussen 1992), one can design an environment so that the elements of the physical and social environment afford stress reduction and enhanced intrinsic motivation. An affordance is a relational property that does not exist independent of the individual and the environment. Affordances therefore share conceptual elements of person–environment transactions that drive emotion and stress. They differ from each other in that the classical definition of affordance often describes it as a physical property of the environment (Gibson 1966, 1979), although more recent thinking suggests that no specific physical element connotes affordance. Thus, one cannot define either concept by isolating either the individual or the context (see Reed 1996).

Motivational affordances may be conceived as elements of the work environment that facilitate and nurture intrinsic motivation. The key for design is to identify *motivational invariants*, or environmental factors that consistently determine an individual’s level of intrinsic motivation across contexts. There are some aspects of work that have been identified as important for facilitating intrinsic motivation and would thus be considered motivational invariants. For instance, providing feedback that is perceived as controlling rather than informative tends to undermine a sense of autonomy and competence and thereby reduces intrinsic motivation (Deci, Ryan, and Koestner 1999). Careful analyses of the motivational affordance structure permit design of tasks that are more likely to be enjoyable by rendering the tools convivial (Illich 1973) and thereby facilitating human–machine synergy (see Hancock 1997).

4.5 DIRECTIONS FOR FUTURE HUMAN–COMPUTER INTERACTION RESEARCH IN RELATION TO WORKLOAD AND STRESS

In this section, we identify directions for future research. These include a better understanding of resources and quantifying task dimensions defined in the Hancock and Warm (1989) model. Progress here likely reduces to the thorny problem of quantifying human information processing (see Hancock, Szalma, and Oron-Gilad 2005). Further, we discuss the need for research on performance–workload associations and dissociations, and the evident need for programmatic investigation of the role of individual differences in performance, workload, and stress.

The Hancock and Warm (1989) model of stress explicitly identifies task dimensions that influence stress state and behavioral adaptability. However, the metrics for these dimensions, and how specific task characteristics map to them, have yet to be fully articulated. Thus, future research should aim to examine how different task components relate to performance and subjective and physiological states. Development of a quantitative model of task characteristics will permit the derivation of vectors for the prediction of adaptability under stress. Cognitive neuroscience and neuroergonomics in particular offer a very promising approach to such understanding. An additional step in this direction, however, will be facilitated by improved quantitative models of how human beings process information (Hancock, Szalma, and Oron-Gilad 2005).

4.5.1 UNDERSTANDING MENTAL RESOURCES

One of the challenges for quantifying human information processing is that there is little understanding or consensus regarding the capacities that “process” the information. A central concept in energetic models of human performance is mental resources. Resource theory replaced arousal and serves as an intervening variable to explain the relations between task demand and performance. However, a continual problem for the resource concept is to operationally define what resources actually are. Most early treatments of resources used that term metaphorically (Navon and Gopher 1979; Wickens 1980, 1984), and the failure to specify what resources have led some to challenge the utility of the concept (Navon 1984). As resource theory is a central concept in theories of stress and represents one of the most important issues to be resolved in future research on stress and performance, we now turn to the definitional concerns associated with the resource construct and imperatives for future research to refine the concept.

4.5.1.1 Resource Metaphors

Two general categories of resource metaphors may be identified as structural metaphors and energetic metaphors. One of the earliest conceptualizations of resource capacity used a computer-based metaphor (Moray 1967). Thus, cognitive capacity was viewed as being analogous to the random

access memory and processing chip of a computer, consisting of information-processing “units” that can be deployed for task performance. However, the structural metaphor has been applied more to theories of working memory than to attention and resource theory.* Most early resource theories, including Kahneman’s (1973) original view and modifications by Norman and Bobrow (1975), Navon and Gopher (1979), and Wickens (1980, 1984), applied energetic metaphors to resources. These perspectives conceptualized resources as commodities or as pools of energy to be “spent” on task performance. In general, energetic approaches tend to employ either economic or thermodynamic/hydraulic metaphors. The economic model is reflected in the description of resources in terms of supply and demand: Performance on one or more tasks suffers when resource demands of the tasks exceed available supply. Presumably the total amount of this supply fluctuates with the state of the individual, and the “assets” diminish with increases in the intensity or duration of stress. Although Kahneman’s (1973) original conception allows for dynamic variation of available resource capacity, most early models assumed a fixed amount of resources (see Navon and Gopher 1979). In thermodynamic analogies, resources comprise a fuel that is consumed or a tank of liquid to be divided among several tasks, and under stressful conditions the amount of resources available is insufficient to meet demand and thus performance suffers. There is no consensus as to the capacity or flexibility of such a tank or whether there are numerous malleable tanks that only store modality-specific resources (Young and Stanton 2002). In discussing his version of resource theory, Wickens (1984) warned that the hydraulic metaphor should not be taken too literally, but most subsequent descriptions of resources have employed visual representations of resources in just this form (i.e., a tank of liquid). Similarly, many discussions of resource availability and expenditure adopt the economic language of supply and demand, and Navon and Gopher (1979) explicitly adopted principles of microeconomics in developing their approach. An additional problem faced by resource theory is that in most cases (e.g., Navon and Gopher 1979; Wickens 1980, 1984) the structural and energetic metaphors are treated interchangeably, a further testament to the ambiguity of the construct.

A problem with using nonbiological metaphors to represent biological systems is that such models often fail to capture the complexity and the unique dynamic characteristics (e.g., adaptive responses) of living systems. For instance, a hydraulic model of resources links the activity of a tank of liquid, governed by thermodynamic principles, to the action of arousal mechanisms or energy reserves that are allocated for task performance. However, a thermodynamic description of the physiological processes underlying resources is at a level of explanation that may not adequately describe the psychological processes that govern performance. Thermodynamic principles can be applied to the chemical processes that occur within and between neurons, but they

may be less useful in describing the behavior of large networks of neurons.[†] Similarly, economic metaphors of supply and demand may not adequately capture the relation between cognitive architecture and energy allocated for their function. Economic models of resources define them as commodities to be spent on one or more activities and they assume an isomorphism between human cognitive activity and economic activity, an assumption that may not be tenable. Indeed, Navon and Gopher (1979) admitted that their “static” economic metaphor for multiple resources may need to be replaced by a dynamic one that includes temporal factors (e.g., serial versus parallel processing; activity of one processing unit being contingent on the output of another). Such concerns over the metaphors used to describe resources are hardly new (Navon 1984; Wickens 1984); but the use of metaphors has become so ingrained in the general scientific thinking about resources and human performance that reevaluation of metaphors is more than warranted, it should be mandated. A regulatory model based on neurophysiological chemistry may serve as a better metaphor (and, in future, may serve to describe the actual nature of the resources themselves to the extent that they can be established) to describe the role of resources in human cognition and performance. However, even a physiology-based theory of resources must be tempered by the problems inherent in reducing psychological processes to physiological activity.

4.5.1.2 Function of Resources

Another problem faced by resource theory is the absence of a precise description of how resources control different forms of information processing. Do resources determine the energy allocated to an information processor (Kahneman 1973), do they provide the space within which the processing structure works (Moray 1967), or does the processor draw on the resources as needed (and as made available)? In the third case, the cognitive architecture would drive energy consumption and allocation, but the locus of control for the division of resources remains unspecified in any case. Presumably an “executive” function that either coordinates information processors drawing on different pools of resources or decides how resources will be allocated must itself consume resources, in terms of both energy required for decision making and mental “space” or structure required. Hence, resource theory does not solve the homunculus problem for theories of attention nor does it adequately describe the resource allocation strategies behind the performance of information-processing tasks.

4.5.1.3 Empirical Tests of the Model

Navon and Gopher (1979) commented on the problem of empirically distinguishing declines in performance due to insufficient supply from those resulting from increases in

* This is a curious historical development, since these relatively separate areas of research converge on the same psychological processes.

[†] The argument here is not that neural structures are not constrained by the laws of thermodynamics—clearly they are—but that thermodynamic principles implied by the metaphor are not sufficient for the development of a complete description of resources and their relation to cognitive activity.

demand. They asked, "When the performance of a task deteriorates, is it because the task now gets fewer resources or because it now requires more?" (Navon and Gopher 1979, p. 243). Navon and Gopher (1979) characterized the problem as distinguishing between changes in resources and changes in the subject-task parameters that constrain resource utilization. They offered two approaches to avoid this conundrum: One approach is to define the fixed constraints of a task and then observe how the information-processing system manages the processes within those constraints. The degree of freedom of the system, in this view, is the pool of resources available, in which the term resource is interpreted broadly to include quality of information, number of extracted features, or visual resolution. The subject-task parameters define what is imposed on the system (the demands) and the resources refer to what the system does in response to the demands (allocation of processing units). From this perspective, resources can be manipulated by the information-processing system within the constraints set by the subject-task parameters. A second approach is to distinguish the kind of control the system exerts on resources between control on the use of processing devices (what we have called structure) and control of the properties of inputs that go into these devices. The devices are *processing resources*. The other kind of control is exerted on *input resources*, which represents the flexibility a person has for determining which inputs are operated on, as determined by subject-task parameters. Processing resources are limited by the capacities of information processors, whereas input resources are limited by subject-task parameters (and allocation strategies that determine which information the operator attends to). Presumably the individual has some control over the allocation strategy, in terms of the processing resources devoted to a task, although these can also be driven by task demands (e.g., a spatial task requires spatial processing units). Navon and Gopher did not advocate either approach but presented both approaches as alternatives for further investigation. The implication for examining the resource model of stress is that one must manipulate both the subject-task parameters (e.g., by varying the psychophysical properties of the stimulus, manipulating the state of the observer, or varying the kind of information processing demanded by the task) and the allocation strategies used by the operator (the input resources, e.g., payoff matrices, task instructions). This would provide information regarding how specific stressors impair specific information-processing units and how they change a user's resource allocation strategies in the presence of stress that is continuously imposed on operators of complex computer-based systems.

In a later article, Navon (1984) moved to a position that is less favorable toward resources than his earlier approach, asserting that predictions derived by resource theory could be made, and results explained, without appealing to the resource concept at all (see also Rugg 1986). One could instead interpret effects in terms of the outputs of information processors. Most manipulations, such as task difficulty (which in Navon's view influences the efficiency of a unit of

resources) or complexity (which affects the load, or the number of operations required), influence the demand for processing, with supply having no impact on their interaction. However, this approach assumes a clear distinction between outputs of a processing system and the concept of a resource, and Navon's (1984) notion of specific processors seems blurred with the notion of a resource, as both are utilized for task performance. Nevertheless, his critique regarding the vagueness of the resource concept is relevant, and Navon did argue that if resources are viewed as an intervening variable rather than a hypothetical construct the concept has utility.

4.5.1.4 Structural Mechanisms

If different kinds of information processing draw on different kinds of resources, in terms of the information processors engaged in a task, stressors may have characteristic effects on each resource. In addition, as Navon and Gopher (1979) have noted, an aspect of resource utilization is the efficiency of each resource unit. It may be that stress degrades the efficiency of information-processing units, independent of energy level or allocation strategy (cf. Eysenck and Calvo 1992). Investigation of such effects could be accomplished by transitioning between tasks requiring different kinds of information processing and determining if the effects of stress on one structure impacts the efficiency of a second structure.

The quality of resources can vary in terms of not only the kind of information-processing unit engaged but also the kind of task required. Following Rasmussen's (1983) classification system for behavior as a heuristic for design some tasks require knowledge-based processing, in which the operator must consciously rely on his or her mental model of the system in order to achieve successful performance. Other tasks fall under the category of rule-based behavior, in which a set of rules or procedures defines successful task response. The third category is skill-based behavior, in which a task is performed with a high degree of automaticity. Presumably each kind of task requires different amounts of resources, but they may also represent qualitatively different forms of resource utilization. In other words, these tasks may differ in the efficiency of a unit of resources as well as in effort allocation strategies. As task performance moves from knowledge- to rule- to skill-based processing (e.g., with training), the cognitive architecture may change such that fewer information-processing units are required and those that are engaged in the performance become more efficient. Moreover, the way in which each of these systems degrade with time under stress may be systematic with the more fragile knowledge-based processing degrading first, followed by rule-based processing, and skill-based processing degrading last (at this point, one may begin to see breakdown of not only psychological processes but also physiological ones; see Hancock and Warm 1989). This degradation may follow a hysteresis function such that a precipitous decline in performance occurs as the operator's resource capacity is reduced below a minimum threshold for performance. Moreover, these processes may recover in an inverse form with skill-based processing

recovering first, followed by rule-based and knowledge-based processing.

Note that it may be difficult to distinguish “pure” knowledge-based processing from rule- or skill-based activity. An alternative formulation is the distinction between controlled and automatic processing (Schneider and Shiffrin 1977). Although originally conceived as categories, it is likely that individuals engaged in real-world tasks utilize both automatic and controlled processing for different aspects of performance and that for a given task there are levels of automaticity possible. Treating skills as a continuum rather than as discrete categories may be a more theoretically useful framework for quantifying resources and information processing, and thereby elucidating the effects of stress on performance.

4.5.1.5 Energetic Mechanisms

To investigate the energetic aspects of resources, one must manipulate environment-based perturbations, in the form of external stressors (noise, heat) and task demands, to systematically affect inflow versus outflow of energy. Presumably inflow is controlled by arousal levels, physiological energy reserves, and effort. One could examine performance under manipulations of energetic resources under dual-task performance (e.g., what happens to the performance on two tasks under conditions of sleep deprivation or caffeine consumption?). For example, the steady state can be perturbed by increasing (e.g., caffeine) or decreasing (e.g., sleep deprivation) energy while systematically varying the demands for the two tasks.

4.5.1.6 Structure and Energy

Another empirical challenge is to distinguish resources as structure from resources as energy. Given the definitional problems associated with the resource concept, it is not clear whether performance declines because of reduction in energy level or degradation of structures (i.e., failures or declines in the efficiency of processing units) or a combination of both. If structure and energy are distinct elements of resources it is hypothetically possible to manipulate one while holding the other constant, although the validity of this assumption is questionable. Is it possible to manipulate specific forms of information processing under constant energy level? Is it possible to manipulate energy level independent of which cognitive processes are utilized? If the decline in available resources is, at least in part, due to the degradation of particular information-processing units, then transferring to a task requiring the same processor should lead to worse performance than transferring to one that is different (cf. Wickens 1980, 1984). For instance, if a person engages in a task requiring verbal working memory while under stress and then transitions to a task requiring spatial discrimination, performance on the latter should depend only on energetic factors and not on structural ones. Note, however, that in this case the effects of different mental capacities would be confounded with the effects of novelty and motivation on performance.

4.5.1.7 Application of Neuroergonomics

The burgeoning field of neuroergonomics seeks to identify the neural bases of psychological processes involved in real-world human–technology interaction (Parasuraman 2003). As we state in Section 4.5 (Hancock and Szalma 2007), recent advances in neuroergonomics promise to identify cognitive processes and their link to neurological processes. For instance, the cognitive process of emotion regulation has been linked to genetic variation in the regulation of neuronal processes. Neuroscientists have isolated a particular genetic polymorphism, 5-HTTLPR, which moderates the level of an individual’s emotional reactivity. The extent of emotional reaction influences the type and amount of mental resources mobilized for its regulation, which potentially signifies far-reaching effects for the entire HCI (Pezawas et al. 2005). Neuroergonomic research may therefore permit a more robust and quantitative definition of resources, although we caution that a simple reductionist approach is not likely to be fruitful as might initially be conceived (see Hancock and Szalma 2003b). In addition, the stress concept itself rests in part on more precise definitions of resources (Hancock and Szalma 2007). Thus, resolution of the resource issue with respect to cognitive processing and task performance would also clarify the workload and stress concepts. We therefore view neuroergonomics as one promising avenue for future research to refine the workload/stress and resource concepts.

4.5.2 DEVELOPMENT OF ADAPTATION UNDER THE STRESS MODEL

Effective adaptation has always been key when performing in demanding environments. This section addresses how adaptation comes about.

4.5.2.1 Quantify the Task Dimensions

A major challenge for the Hancock and Warm (1989) model is the quantification of the base axes representing task dimensions. Specification of these dimensions is necessary if the vector representation postulated by Hancock and Warm is to be developed and if the resource construct is to be more precisely defined and quantified. However, task taxonomies that are general across domains present a theoretical challenge, because they require an understanding and quantification of how individuals process information along the spatial and temporal task dimensions, and how these change under stressful conditions. Quantification of information processing, and subsequent quantification of the base axes in the Hancock and Warm (1989) model, will permit the formalization of the vector representation of adaptive state under stress (see Figure 4.4).

4.5.2.2 Attentional Narrowing

Recall that Hancock and Weaver (2002) argued that the distortions of spatial and temporal perception have a common attentional mechanism. Two implications of this assertion are as follows: (1) Events (internal or external) that distort one dimension will distort the other and (2) these distortions

are unlikely to be orthogonal. With very few exceptions, little research has addressed the possibility of an interaction between distortions of spatial and temporal perceptions in stressful situations on operator performance. Preliminary evidence suggests that these two dimensions may in fact not share a simple, common mechanism (Ross et al. 2003; Thropp, Szalma, and Hancock 2004), although further research is needed to confirm this finding. An additional important issue for empirical research is whether we are dealing with “time-in-memory” or “time-in-passing” (and to some extent, space-in-memory vs. space-in-passing). Thus, the way in which perceptions of space and time interact to influence operator state will depend on how temporal perceptions (and spatial perception, for that matter) are measured.

A possible explanation for perceptual distortions under conditions of heavy workload and stress concerns the failure to switch tasks when appropriate. Switching failures may be responsible for the observation in secondary task methodology that some participants have difficulty dividing their time between tasks as instructed (e.g., 70% to the primary task and 30% to the secondary task). This difficulty may result from the participant’s inability to accurately judge how long he or she has attended to each task during a given time period. The degree to which distortions in the perception of space–time are related to impairments in task switching under stressful conditions and the degree to which these distortions are related to attention allocation strategies in a secondary task paradigm are questions for empirical resolution.

4.5.2.3 Stressor Characteristics

Even if space and time do possess a simple, common mechanism, it may be that specific stressors do not affect spatial and temporal perceptions in the same way. For instance, heat and noise may distort perception of both space and time but not to the same degree or in the same fashion. It is important to note that spatial and temporal distortions may *themselves* be appraised as stressful, as they might interfere with the information-processing requirements of a task. Consequently, some kinds of information processing might be more vulnerable to one or the other kind of perceptual distortion. Clearly, performance on tasks requiring spatial abilities, such as mental rotation, could suffer as a result of spatial distortion, whereas they might be unaffected (or, in some cases, facilitated) by temporal distortion. Other tasks, such as tasks that rely heavily on working memory or mathematical ability, or tasks requiring target detection, could each show different patterns of change in response to space–time distortion.

4.5.2.4 Potential Benefits of Space–Time Distortion

Under certain conditions, the narrowing of spatial attention can benefit performance through the elimination of irrelevant cues. The precise conditions under which this occurs, however, remains unclear. In addition, it is important to identify the circumstances under which time distortion might actually prove beneficial. Here, operators perceive that they have *additional* time to complete the task at hand (Hancock and Weaver 2005). This has great benefit in task performance

situations where attentional narrowing is less likely to have deleterious effects. At this point in time, this is an empirical question that might be amenable to controlled testing.

4.5.2.5 Changes in Adaptation: The Roles of Time and Intensity

The degree to which a task or the physical and social environment imposes stress is moderated by the characteristics of the stimuli as well as the context in which events occur. However, two factors that seem to ubiquitously influence how much stress impairs adaptation are (appraised) intensity of the stressor and duration of exposure. We have reported meta-analytic evidence that these two factors jointly impact task performance across different orders of tasks (e.g., vigilance tasks, problem solving, tracking; see Hancock, Ross, and Szalma 2007; Szalma and Hancock 2011; Szalma, Hancock, and Quinn 2008). Duration is further implicated in information processing itself and may be a central organizing principle for information processing in the brain. Duration and intensity of environmental stimulation can likewise influence emotional reactions and consequently which emotion regulation or stress-coping strategy an individual opts to employ (Gross 1998, 2002). Empirical research is, however, still needed to explore programmatically the interactive effects of these variables across multiple forms of information processing.

4.5.3 UNDERSTANDING PERFORMANCE–WORKLOAD ASSOCIATIONS/DISSOCIATIONS

Factors that prompt associations or dissociations are herein discussed as well as their contribution to perceived workload.

4.5.3.1 Task Factors

Although Hancock (1996) and Yeh and Wickens (1988) have articulated the patterns of performance–workload relations and how these are diagnostic with respect to processing requirements, little systematic effort has been spent on further investigating these associations/dissociations. The primary question is what factors drive dissociations and insensitivities when they occur. For instance, for vigilance tasks mostly associations are observed, whereas for other tasks, such as those with high working memory demand, dissociations are more common (Yeh and Wickens 1988). Enhanced understanding of these relations would inform the Hancock and Warm (1989) model by permitting specification of the conditions under which individuals pass over the thresholds of failure at each level of person–environment transaction/adaptation.

4.5.3.2 Multidimensionality of Workload

To date, consideration of performance–workload dissociations has been primarily concerned with global measures of perceived workload. However, there is clear evidence that perceived workload is in fact multidimensional. For instance, vigilance tasks are characterized by high levels of mental demand and frustration (Warm, Dember, and Hancock 1996).

It is likely that the pattern of performance–workload links is different not only for different orders of performance (different tasks) but also for different dimensions of workload. One approach to addressing this question would be to systematically manipulate combinations of these two variables. For instance, if we consider performance in terms of detection sensitivity, memory accuracy, and speed of response, and the dimensions of workload defined by the National Aeronautics and Space Administration (NASA) Task Load Index (Hart and Staveland 1988), we could discuss how variations in memory load or discrimination difficulty link to each subscale.

4.6 INDIVIDUAL DIFFERENCES IN PERFORMANCE, WORKLOAD, AND STRESS

In previous work, we have reviewed the relations between individual differences in state and trait and efforts to quantify human information processing (Szalma and Hancock 2005). In this section, we address how individual differences (state and trait) are related to stress and coping.

4.6.1 TRAIT DIFFERENCES

Individual differences research has been a neglected area in human factors and experimental psychology. Much of the early work on individual differences was done by researchers who were unconcerned with human–technology interactions to the extent that a bifurcation between two kinds of psychology occurred (Cronbach 1957). There is evidence, however, that affective traits influence information processing and performance. Thus, extraversion is associated not only with superior performance in working memory tasks and divided attention but also with poorer sustained attention (however, see Koelega 1992). Trait anxiety is associated with poor performance, although results vary across task types and contexts (Matthews, Deary, and Whiteman 2003; Szalma 2008). A possible next step for such research is to systematically vary task elements, as discussed previously (in Section 4.3.1) in the context of the Hancock and Warm model, and test hypotheses regarding how trait anxiety relates to specific task components. The theoretical challenge for such an undertaking is that it requires a good taxonomic scheme for tasks as well as a well-articulated theory of traits and performance. However, trait theories have neglected specific task performance, focusing instead on global measures (e.g., see Barrick, Mount, and Judge 2001), and there is a lack of a comprehensive theory that accounts for trait–performance relations (Matthews, Deary, and Whiteman 2003). Most current theories are more like frameworks that do not provide specific mechanisms for how personality impacts cognition and performance (e.g., see McCrae and Costa 1999). Although Eysenck (1967) proposed a theory of personality based on arousal and activation, which has found some support (Matthews and Gilliland 1999), there is also evidence to the end that arousal and task difficulty fail to interact as predicted (Matthews 1992). Eysenck’s (1967) theory was also

weakened by the general problems associated with arousal theory accounts for stress effects (Hockey 1984). An alternative formulation is that of Gray (1991) who argued for two systems, (1) one responding to reward signals and (2) one to punishment. The behavioral activation system (BAS) is associated with positive affect, whereas the behavioral inhibition system (BIS) is associated with negative affect. In a review and some comparisons of the Eysenck and Gray theories, Matthews and Gilliland (1999) partially supported both the theories but concluded that Gray’s BAS/BIS distinction provides a superior match to positive and negative affect relative to Eysenck’s arousal dimensions. Further, the BAS/BIS distinction accords with theories of approach/avoidance motivation (e.g., Elliot and Covington 2001). Indeed, intraindividual approaches to investigating the complex interplay between stress, coping, and performance outcomes are hailed as the most promising methodology (Folkman et al. 1986). There are also theories that focus on a particular trait such as extraversion (Humphreys and Revelle 1984) or trait anxiety (Eysenck and Calvo 1992). Although useful, such specific theories do not encompass other traits or interactions among traits. Such interactive effects can influence cognitive performance and perceived stress and workload (Szalma et al. 2005). These interactions should be further studied with an eye to linking them to information-processing theories.

4.6.2 AFFECTIVE STATE DIFFERENCES

It is intuitive that stress would induce more negative affective states and that traits would influence performance via an effect on states. For instance, one would expect that trait anxiety would influence performance because high trait anxious individuals experience state anxiety more frequently than those low on that trait. Although such mediation effects are observed, there is also evidence that for certain processes, such as hypervigilance to threat, trait anxiety is a better predictor of performance than state anxiety (Eysenck 1992). In terms of appraisal theory, traits may influence the form and content of appraisals, as well as the coping skills an individual can deploy to deal with stress. With respect to adaptation, it is likely that individual differences in both trait and state will influence adaptation, both behavioral and physiological, by affecting the “width” of the plateau of effective adaptation at a given level and by changing the slope of decline in adaptation when the adaptation threshold is reached. That is, higher skill levels protect from declines in adaptive function by increasing the threshold for failure at a given level (i.e., comfort, performance, physiological response). The modification of the Hancock and Warm (1989) model illustrating these individual differences in effects is shown in Figure 4.6 (and see Szalma 2008). Multiple frameworks of state dimensions exist, but most focus on either two (e.g., Thayer 1989; Watson and Tellegen 1985) or three (Matthews et al. 1999, 2002) frameworks. In the context of task performance, Matthews and his colleagues have identified three broad state dimensions reflecting the cognitive, affective, and motivational aspects of an individual’s current psychological state.




FIGURE 4.6 Modification of the adaptability model shown in Figure 4.3. The adaptability model of Hancock and Warm (1989) shown in Figure 4.3 has been modified to illustrate how individual differences may influence stress and adaptation. It is likely that cognitive and affective traits influence the width of both the comfort and performance zones (i.e., the thresholds for declines in adaptation) as well as the rate of decline in adaptability when a threshold is crossed. For instance, individuals high in trait anxiety likely have a narrower plateau of stability and therefore manifest lower thresholds for discomfort and performance degradation than individuals low on that trait. Further, the rate of decline in adaptation may increase as a function of trait anxiety.

These dimensions are “worry,” which reflects the cognitive dimension of stress, and “task engagement” and “distress,” which reflect the affective, cognitive, and motivational components of state. Specifically, a high level of distress is indicative of overload in processing capacity, and task engagement reflects a theme of commitment to effort (Matthews et al. 2002). Matthews and his colleagues have demonstrated that changes in task demand influence the pattern of stress state. Should affective state fail to detrimentally influence task performance itself, it can still critically impact consequential levels of “psychophysiological activation, strain, and fatigue aftereffects” (Robert and Hockey 1997, p. 73). It is therefore important to incorporate assessment of operator state into the interface design process so that the interaction with technology fosters task engagement and minimizes distress and worry.

4.6.2.1 Attentional Narrowing and Adaptive Response

As with other aspects of perception, there are individual differences in the perception of space and time (Hancock and Weaver 2005; Wachtel 1967). Further, because the subjective experience of stress is often multidimensional, it may be that although two individuals are subjectively stressed by the same situation, their stress profiles differ. Affective states can likewise influence the extent of attentional allocation. Affective states high in motivational intensity, either pleasant or unpleasant, lead to a narrowing of attentional focus, whereas affective states low in motivational intensity, again

regardless of valence, cause attentional broadening (Gable and Harmon-Jones 2009). Individuals are also likely to differ in the strategies they employ to cope with the distortions of space–time and emotional flux they experience while in a stressful environment, and these coping differences, if they exist, might depend on the quality (e.g., noise, heat, low signal salience) and source (e.g., the environment, the task) of the stress and the personality traits of an individual.

4.6.2.2 Hedonomics and Individual Differences

In addition to application of individual differences research to the development of training or selection procedures, individual difference variables can be used to promote hedonomic approaches to design and facilitate interface design. Thus, if the traits that influence the subjective experience of an interaction with technology are identified, that interface can then be configured to meet the preferences and the trait/state profile of an individual user and promote positive affective states. However, for such efforts to succeed, the relations among traits and cognitive, perceptual, and motor performance need to be established via theory-guided empirical research.

4.7 IMPLICATIONS OF STRESS FOR RESEARCHERS AND PRACTITIONERS

For both research and design applications, extant research on stress and performance indicates that assessments of workload and affective state are important for a more complete understanding of HCI. Such assessments can aid in

identifying which components of an interface or task are appraised as stressful, and thus interfaces can be designed to mitigate their negative effects. For instance, research is needed to establish which task parameters control the patterns of performance–workload associations and dissociations and how these change dynamically as a function of time on task. The Hancock and Warm (1989) model of stress established general task dimensions (space–time) that influence stress state and behavioral adaptability, but the metrics for these dimensions remain elusive. This problem results from the central issue regarding how to quantify human information processing (Hancock, Szalma, and Oron-Gilad 2005) and define “mental resources” more precisely (Hancock and Szalma 2007). Efforts to resolve these definitional problems would improve stress theory and its application to interface design. Future research should therefore examine the relations between task dimensions and user characteristics and how these change over time and under high-stress conditions.

In addition to changing the task, there are other techniques that can be applied to the design of human–computer interfaces for use in stressful environments. These include skill development (e.g., Hancock 1986), use of configural displays (Hancock and Szalma 2003a; Wickens 1996), as well as use of technologies employing adaptive automation and decision aids (Hancock and Chignell 1987). With respect to skill development in particular, an area in need of research is how individuals who develop expertise in a task also learn how to cope with stress while performing the task. In order to understand how individuals accomplish this, one is required to understand in depth how different forms of stress influence different forms of information processing. Intuitively, automation and decision aids seem to be key tools for relieving stress during task performance. Although experiments have yielded some promising results (Funke et al. 2007), further research is necessary to determine this supposition under different kinds of stress as such technologies could merely serve to divert attentional resources away from the task.

It is also important for both researchers and practitioners to consider the characteristics of a user and how these characteristics interact with the task or interface to influence performance.

An understanding of how individual differences influence HCI can facilitate the development of tailored training regimens as well as interfaces that can more effectively adapt to the user. Systems that can respond to changes in operator affective state can achieve the desired human–machine synergy in HCI (cf. Hancock 2009). Realizing these goals, however, will require adequate theory development and subsequent empirical research to determine the nature of the relations among the person and environmental variables. It is particularly important to design interfaces that permit autonomous motivation (Deci and Ryan 2000) and to understand how operators of computer-based systems can internalize extrinsic motivation as they gain experience on the task (Gagne and Deci 2005). We suggest here that researchers and designers identify the motivational affordances in the task environment and utilize them to enhance the experience of

HCI and improve overall system performance under stress. Motivational affordances will be elements of the work environment that facilitate and nurture intrinsic motivation. Particularly important for design will be the identification of motivational invariants, which are those environmental factors that consistently determine an individual’s level of intrinsic (or extrinsic) motivation across contexts. Careful analyses of the motivational affordance structure will permit design of tasks that are more likely to be enjoyable by rendering the tools convivial (Illich 1973) and thereby facilitating human–machine synergy (see Hancock 1997).

4.8 SUMMARY AND CONCLUSIONS

In this chapter, we review theories of stress and performance and their relevance to human–technology interaction. We also show that despite being developed in separate research traditions, workload and stress can be viewed as different perspectives on the same fundamental problem. We outline general principles for stress mitigation and discuss issues that require further research. Of particular importance are establishing sound measures of information processing and mental resource expenditure as well as articulating the relevant task dimensions and how they trigger self-regulatory mechanisms, specifically emotion (stress) regulation techniques. Given that stress can be understood only in relation to the transaction between an individual and the environment, it is crucial to establish how trait and state characteristics of the individual influence their appraisals. Finally, it is important in practical applications to treat stress at multiple levels, ranging from the physiological to the organizational sources of adverse performance effects. Different emotion regulation strategies attempt to mitigate stress at these various levels; which techniques are chosen by an operator to utilize can significantly influence the success of an HCI. Traditional attempts to treat stress problems unidimensionally will continue to fail until the person, task, and the physical and social/organizational environment are treated by analysis as a coherent system. Researchers and practitioners in HCI should therefore expand their efforts beyond the design of displays and controls of interfaces and include assessments of person-related factors that influence performance as well as the design of the physical and social environment in which an HCI occurs.

REFERENCES

- Ackerman, P. L., and E. D. Heggstad. 1997. Intelligence, personality, and interests: Evidence for overlapping traits. *Psychol Bull* 121:219–45.
- Aspinwall, L. G., L. Richter, and R. R. Hoffman. 2002. Understanding how optimism works: An examination of optimists’ adaptive moderation of belief and behavior. In *Optimism and Pessimism: Implications for Theory, Research, and Practice*, ed. E. C. Chang, 217–38. Washington, DC: American Psychological Association.
- Bandura, A. 1997. *Self-Efficacy: The Exercise of Control*. New York: W.H. Freeman and Company.

- Barrick, M. R., M. K. Mount, and T. A. Judge. 2001. Personality and performance at the beginning of the new millennium: What do we know and where do we go next? *Pers Perform* 9:9–29.
- Bennett, K. B., and J. M. Flach. 1992. Graphical displays: Implications for divided attention, focused attention, and problem solving. *Hum Factors* 34:513–52.
- Bless, H. 2001. Mood and the use of general knowledge structures. In *Theories of Mood and Cognition: A User's Guidebook*, ed. L. L. Martin and G. L. Clore, 9–26. Mahwah, NJ: Erlbaum.
- Brown, J. M., and E. A. Campbell. 1994. *Stress and Policing: Sources and Strategies*. Chichester, UK: Wiley.
- Bursill, A. E. 1958. The restriction of peripheral vision during exposure to hot and humid conditions. *Q J Exp Psychol* 10:113–29.
- Carver, C. S., and M. F. Scheier. 1998. *On the Self-Regulation of Behavior*. New York: Cambridge University Press.
- Conway, G. E., J. L. Szalma, and P. A. Hancock. 2007. A quantitative meta-analytic examination of whole-body vibration effects on human performance. *Ergonomics* 50:228–45.
- Cornsweet, D. M. 1969. Use of cues in the visual periphery under conditions of arousal. *J Exp Psychol* 80:14–8.
- Cronbach, L. J. 1957. The two disciplines of scientific psychology. *Am Psychol* 12:671–84.
- Csikszentmihalyi, M. 1990. *Flow: The Psychology of Optimal Experience*. New York: Harper & Row.
- Deci, E. L., and R. M. Ryan. 1985. *Intrinsic Motivation and Self-Determination in Human Behavior*. New York: Plenum Press.
- Deci, E. L., and R. M. Ryan. 2000. The “what” and “why” of goal pursuits: Human needs and the self-determination of behavior. *Psychol Inq* 11:227–68.
- Deci, E. L., R. M. Ryan, and R. Koestner. 1999. A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychol Bull* 125:627–68.
- Dirkin, G. R., and P. A. Hancock. 1984. Attentional narrowing to the visual periphery under temporal and acoustic stress. *Aviat Space Environ Med* 55:457.
- Dirkin, G. R., and P. A. Hancock. 1985. An attentional view of narrowing: The effect of noise and signal bias on discrimination in the peripheral visual field. In *Ergonomics International 85: Proceedings of the Ninth Congress of the International Ergonomics Association*, ed. I. D. Brown, R. Goldsmith, K. Coombes, and M. A. Sinclair, 751–3. Bournemouth, England, London: Taylor & Francis.
- Driskell, J. E., J. H. Johnston, and E. Salas. 2001. Does stress training generalize to novel settings? *Hum Factors* 43:99–110.
- Easterbrook, J. A. 1959. The effect of emotion on cue utilization and the organization of behavior. *Psychol Rev* 66:183–201.
- Elliot, A. J., and M. V. Covington. 2001. Approach and avoidance motivation. *Educ Psychol Rev* 13:73–92.
- Ericsson, K. A. 2006. The influence of experience and deliberate practice on the development of superior expert performance. In *The Cambridge Handbook of Expertise and Expert Performance*, ed. K. A. Ericsson, N. Charness, R. R. Hoffman, and P. J. Feltovich, 683–703. Cambridge: Cambridge University Press.
- Eysenck, H. J. 1967. *The Biological Basis of Personality*. Springfield, IL: Charles C. Thomas.
- Eysenck, M. W. 1992. *Anxiety: The Cognitive Perspective*. Hillsdale, NJ: Erlbaum.
- Eysenck, M. W., and M. G. Calvo. 1992. Anxiety and performance: The processing efficiency theory. *Cogn Emot* 6:409–34.
- Flach, J., P. A. Hancock, J. K. Caird, and K. Vicente, eds. 1995. *Global Perspectives on the Ecology of Human-Machine systems*, vol. 1. Hillsdale, NJ: Erlbaum.
- Folkman, S., R. S. Lazarus, C. Dunkel-Schetter, A. DeLongis, and R. J. Gruen. 1986. Dynamics of a stressful encounter: Cognitive appraisal, coping, and encounter outcomes. *J Pers Soc Psychol* 50:992–1003.
- Fromm, E. 1976. *To Have or to be?* New York: Harper & Row.
- Funke, G., G. Matthews, J. S. Warm, and A. K. Emo. 2007. Vehicle automation: A remedy for driver stress? *Ergonomics* 50:1302–23.
- Gable, P., and E. Harmon-Jones. 2009. The blues broaden, but the nasty narrows: Attentional consequences of negative affects low and high in motivational intensity. *Psychol Sci* 20:1–5.
- Gagne, M., and E. L. Deci. 2005. Self-determination theory and work motivation. *J Organ Behav* 26:331–62.
- Gibson, J. J. 1966. *The Senses Considered as Perceptual Systems*. Boston: Houghton Mifflin.
- Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Gray, J. A. 1991. Neural systems, emotion, and personality. In *Neurobiology of Learning, Emotion, and Affect*, ed. J. Madden IV, 273–306. New York: Raven Press.
- Gross, J. J. 1998. Antecedent- and response-focused emotion regulation: Divergent consequences for experience, expression, and physiology. *J Pers Soc Psychol* 74:224–37.
- Gross, J. J. 2002. Emotion regulation: Affective, cognitive, and social consequences. *Psychophysiology* 39:281–91.
- Gross, J. J., and R. W. Levenson. 1993. Emotional suppression: Physiology, self-report, and expressive behavior. *J Pers Soc Psychol* 64:970–86.
- Gross, J. J., J. M. Richards, and O. P. John. 2006. Emotion regulation in everyday life. In *Emotion Regulation in Couples and Families: Pathways to Dysfunction and Health*, ed. D. K. Snyder, J. A. Simpson, and J. N. Hughes, 13–35. Washington, DC: American Psychological Association.
- Haga, S. M., P. Kraft, and E. K. Corby. 2009. Emotion regulation: Antecedents and well-being outcomes of cognitive reappraisal and expressive suppression in cross-cultural samples. *J Happiness Stud* 10:271–91.
- Hancock, P. A. 1986. The effect of skill on performance under an environmental stressor. *Aviat Space Environ Med* 57:59–64.
- Hancock, P. A. 1996. Effects of control order, augmented feedback, input device and practice on tracking performance and perceived workload. *Ergonomics* 39:1146–62.
- Hancock, P. A. 1997. *Essays on the Future of Human-Machine Systems*. Eden Prairie, MN: BANTA Information Services Group.
- Hancock, P. A. 2009. *Mind, machine and morality*. Chichester, UK: Ashgate.
- Hancock, P. A. 2010. The battle for time in the brain. In *Time, Limits and Constraints: The Study of Time XIII*, ed. J. A. Parker, P. A. Harris, and C. Steineck. Leiden, the Netherlands: Brill.
- Hancock, G. M., and G. F. Beatty. 2010. The effects of automatic and deliberate emotion regulation on sustained motor force production. *Program of the University of Florida's 2010 Graduate Student Council Interdisciplinary Research Conference*, 39 (Abstract Only). Gainesville, FL: UF's Graduate Student Council.
- Hancock, P. A., and J. K. Caird. 1993. Experimental evaluation of a model of mental workload. *Hum Factors* 35:413–29.
- Hancock, P. A., and M. H. Chignell. 1987. Adaptive control in human-machine systems. In *Human Factors Psychology*, ed. P. A. Hancock, 305–45. Amsterdam: Elsevier.
- Hancock, P. A., and G. R. Dirkin. 1983. Stressor induced attentional narrowing: Implications for design and operation of person-machine systems. *Proc Hum Factors Assoc Can* 16:19–21.

- Hancock, G. M., and P. A. Hancock. 2010. Can technology create instant experts? *Ergonomist* 460:4–5.
- Hancock, P. A., A. A. Pepe, and L. L. Murphy. 2005. Hedonomics: The power of positive and pleasurable ergonomics. *Ergon Des* 13:8–14.
- Hancock, P. A., J. M. Ross, T. Oron-Gilad, and J. L. Szalma. 2005. *The Incorporation of Comprehensive Thermal Stress Effects into IMPRINT* (Tech. Rep. No. 1). Orlando, FL: University of Central Florida, Minds in Technology, Machines in Thought Laboratory.
- Hancock, P. A., J. M. Ross, and J. L. Szalma. 2007. A meta-analysis of performance response under thermal stressors. *Hum Factors* 49:851–77.
- Hancock, P. A., and J. L. Szalma. 2003a. Operator stress and display design. *Ergon Des* 11:13–8.
- Hancock, P. A., and J. L. Szalma. 2003b. The future of neuroergonomics. *Theor Issues Ergon Sci* 4:238–49.
- Hancock, P. A., and J. L. Szalma. 2007. Stress and neuroergonomics. In *Neuroergonomics: The Brain at Work*, ed. R. Parasuraman and M. Rizzo, 195–206. Oxford: Oxford University Press.
- Hancock, P. A., J. L. Szalma, and T. Oron-Gilad. 2005. Time, emotion, and the limits to human information processing. In *Quantifying Human Information Processing*, ed. D. K. McBride and D. Schmorrow, 157–75. Lanham, MD: Lexington Books.
- Hancock, P. A., and J. S. Warm. 1989. A dynamic model of stress and sustained attention. *Hum Factors* 31:519–37.
- Hancock, P. A., and J. L. Weaver. 2005. On time distortion under stress. *Theor Issues Ergon Sci* 6:193–211.
- Harris, W. C., P. A. Hancock, and S. C. Harris. 2005. Information processing changes following extended stress. *Mil Psychol* 17:115–28.
- Hart, S. G., and L. E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Human Mental Workload*, ed. P. A. Hancock and N. Meshkati, 139–83. Amsterdam: Elsevier.
- Hassenzahl, M., and D. Ullrich. 2007. To do or not to do: Differences in user experience and retrospective judgments depending on the presence or absence of instrumental goals. *Interact Comput* 19:429–37.
- Hockey, R. 1984. Varieties of attentional state: The effects of environment. In *Varieties of Attention*, ed. R. Parasuraman and D. R. Davies, 449–83. New York: Academic Press.
- Hockey, G. R. J. 1997. Compensatory control in the regulation of human performance under stress and high workload: A cognitive–energetical framework. *Biol Psychol* 45:73–93.
- Hockey, G. R. J., A. W. K. Gaillard, and M. G. H. Coles, eds. 1986. *Energetic Aspects of Human Information Processing*. The Netherlands: Nijhoff.
- Hockey, R., and P. Hamilton. 1983. The cognitive patterning of stress states. In *Stress and Fatigue in Human Performance*, ed. G. R. J. Hockey, 331–62. Chichester: Wiley.
- Humphreys, M. S., and W. Revelle. 1984. Personality, motivation, and performance: A theory of the relationship between individual differences and information processing. *Psychol Rev* 91:153–84.
- Illich, I. 1973. *Tools for Conviviality*. New York: Harper & Row.
- Johnston, J. H., and J. A. Cannon-Bowers. 1996. Training for stress exposure. In *Stress and Human Performance*, ed. J. E. Driskell and E. Salas, 223–56. Mahwah, NJ: Erlbaum.
- Kahneman, D. 1973. *Attention and Effort*. Englewood Cliffs, NJ: Prentice Hall.
- Kasser, T. 2002. *The High Price of Materialism*. Cambridge, MA: MIT Press.
- Koelega, H. S. 1992. Extraversion and vigilance performance: Thirty years of inconsistencies. *Psychol Bull* 112:239–58.
- Lang, P. J. 1995. The emotion probe: Studies of motivation and attention. *Am Psychol* 50:372–85.
- Lazarus, R. S. 1991. *Emotion and Adaptation*. Oxford: Oxford University Press.
- Lazarus, R. S. 1999. *Stress and Emotion: A New Synthesis*. New York: Springer.
- Lazarus, R. S., and S. Folkman. 1984. *Stress, Appraisal, and Coping*. New York: Springer Verlag.
- Lepper, M. R., and D. I. Cordova. 1992. A desire to be taught: Instructional consequences of intrinsic motivation. *Motiv Emot* 16:187–208.
- Luczak, H., M. Roettling, and L. Schmidt. 2003. Let's talk: Anthropomorphization as means to cope with stress of interacting with technical devices. *Ergonomics* 46:1361–74.
- Matthews, G. 1992. Extraversion. In *Handbook of Human Performance*, vol. 3: *State and Trait*, ed. A. P. Smith and D. M. Jones, 95–126. London: Academic Press.
- Matthews, G. 2001. Levels of transaction: A cognitive science framework for operator stress. In *Stress, Workload, and Fatigue*, ed. P. A. Hancock and P. A. Desmond, 5–33. Mahwah, NJ: Erlbaum.
- Matthews, G., I. J. Deary, and M. C. Whiteman. 2003. *Personality Traits*. 2nd ed. Cambridge: Cambridge University Press.
- Matthews, G., and K. Gilliland. 1999. The personality theories of H. J. Eysenck and J. A. Gray: A comparative review. *Pers Individ Dif* 26:583–626.
- Matthews, G., L. Joyner, K. Gilliland, S. Campbell, S. Falconer, and J. Huggins. 1999. Validation of a comprehensive stress state questionnaire: Towards a state ‘big three’? In *Personality Psychology in Europe*, ed. I. Mervielde, I. J. Deary, F. De Fruyt, and F. Ostendorf, vol. 7, 335–50. Tilburg, The Netherlands: Tilburg University Press.
- Matthews, G., L. Joyner, K. Gilliland, S. Campbell, S. Falconer, J. Huggins, K. Gilliland, R. Grier, and J. S. Warm. 2002. Fundamental dimensions of subjective state in performance settings: Task engagement, distress, and worry. *Emotion* 2:315–40.
- Mauss, I. B., S. A. Bunge, and J. J. Gross. 2007. Automatic emotion regulation. *Soc Personal Psychol Compass* 1:146–67.
- McCrae, R. R., and P. T. Costa. 1986. Personality, coping, and coping effectiveness in an adult sample. *J Pers* 54:385–405.
- McCrae, R. R., and P. T. Costa. 1999. A five-factor theory of personality. In *Handbook of Personality: Theory and Research*, 2nd ed., ed. L. A. Pervin and O. P. John, 139–53. New York: Guilford Press.
- Moray, N. 1967. Where is capacity limited? A survey and a model. *Acta Psychol* 27:84–92.
- Moray, N., ed. 1979. *Mental Workload: Its Theory and Measurement*. New York: Plenum Press.
- Navon, D. 1984. Resources—A theoretical soupstone? *Psychol Rev* 91:216–34.
- Navon, D., and D. Gopher. 1979. On the economy of the human information processing system. *Psychol Rev* 86:214–55.
- Norman, D., and D. Bobrow. 1975. On data-limited and resource-limited processing. *J Cogn Psychol* 7:44–60.
- Ochsner, K. N., and J. J. Gross. 2005. The cognitive control of emotion. *Trends Cogn Sci* 9:242–9.
- O'Donnell, R. D., and F. T. Eggemeier. 1986. Workload assessment methodology. In *Handbook of Human Perception and Performance*, vol. II: *Cognitive Processes and Performance*, ed. K. R. Boff, L. Kaufman, and J. P. Thomas, 42-1–42-49. New York: Wiley.
- Parasuraman, R. 2003. Neuroergonomics: Research and practice. *Theor Issues Ergon Sci* 4:5–20.

- Parasuraman, R., and P. A. Hancock. 2001. Adaptive control of mental workload. In *Stress, Workload, and Fatigue*, ed. P. A. Hancock and P. A. Desmond, 305–20. Mahwah, NJ: Erlbaum.
- Penley, J. A., and J. Tomaka. 2002. Associations among the big five, emotional responses, and coping with acute stress. *Pers Individ Dif* 32:1215–28.
- Pezawas, L., A. Meyer-Lindenberg, E. M. Drabant, B. A. Verchinski, K. E. Munoz, B. S. Kolachana, M. F. Egan, V. S. Mattay, A. R. Hariri, and D. M. Weinberger. 2005. 5-HTTLPR polymorphism impacts human cingulate-amygdala interactions: A genetic susceptibility mechanism for depression. *Nat Neurosci* 8:828–34.
- Pilcher, J. J., E. Nadler, and C. Busch. 2002. Effects of hot and cold temperature exposure on performance: A meta-analytic review. *Ergonomics* 45:682–98.
- Rasmussen, J. 1983. Skills, rules, and knowledge: Signals, signs, and symbols, and other distinctions in human performance models. *IEEE Trans Syst Man Cybern SMC* 13:257–66.
- Reed, E. S. 1996. *Encountering the World: Toward an Ecological Psychology*. Oxford: Oxford University Press.
- Robert, G., and G. R. J. Hockey. 1997. Compensatory control in the regulation of human performance under stress and high workload: A cognitive-energetical framework. *Biol Psychol* 45:73–93.
- Ross, J. M., J. L. Szalma, J. E. Thropp, and P. A. Hancock. 2003. Performance, workload, and stress correlates of temporal and spatial task demands. *Proc Hum Factors Ergon Soc* 47:1712–6.
- Rugg, M. D. 1986. Constraints on cognitive performance: Some problems with and alternatives to resource theory. In *Energetics and Human Information Processing*, ed. G. R. J. Hockey, A. W. K. Gaillard, and M. G. H. Coles, 353–71. Dordrecht: Martinus Nijhoff Publishers.
- Ryan, R. R., and E. L. Deci. 2000. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *Am Psychol* 55:66–78.
- Ryan, R. R., and E. L. Deci. 2001. On happiness and human potentials: A review of research on hedonic and eudaimonic well-being. *Annu Rev Psychol* 52:141–66.
- Saunders, T., J. E. Driskell, J. Johnston, and E. Salas. 1996. The effect of stress inoculation training on anxiety and performance. *J Occup Health Psychol* 1:170–86.
- Scerbo, M. W., F. G. Freeman, and P. J. Mikulka. 2003. A brain-based system for adaptive automation. *Theor Issues Ergon Sci* 4:200–19.
- Scheier, M. F., and C. S. Carver. 1985. Optimism, coping, and health: Assessment and implications of generalized outcome expectancies. *Health Psychol* 4:219–47.
- Scherer, K. R. 1999. Appraisal theory. In *Handbook of Cognition and Emotion*, ed. T. Dalgleish and M. Power, 638–63. New York: Wiley.
- Schneider, W., and R. M. Shiffrin. 1977. Controlled and automatic human information processing I: Detection, search, and attention. *Psychol Rev* 84:1–66.
- Seligman, M. E. P., and M. Csikszentmihalyi. 2000. Positive psychology: An introduction. *Am Psychol* 55:5–14.
- Selye, H. 1976. *The Stress of life. (Revised edition)*. New York: McGraw-Hill.
- Speisman, J. C., R. S. Lazarus, A. Mordkoff, and L. Davison. 1964. Experimental reduction of stress based on ego- defense theory. *J Abnorm Soc Psychol* 68:367–80.
- Szalma, J. L. 2008. Individual differences in stress reaction. In *Performance Under Stress*, ed. P. A. Hancock and J. L. Szalma, 323–57. Aldershot, England: Ashgate.
- Szalma, J. L., and P. A. Hancock. 2005. Individual differences in information processing. In *Quantifying Human Information Processing*, ed. D. K. McBride and D. Schmorow, 177–93. Lanham, MD: Lexington Books.
- Szalma, J. L., and P. A. Hancock. 2011. Noise effects on human performance: A meta-analytic synthesis. *Psychological Bulletin* 137:682–707.
- Szalma, J. L., P. A. Hancock, and S. Quinn. 2008. A meta-analysis of the effect of time pressure of human performance. *Hum Factors Ergon Soc Annu Meet Proc* 52:1513–6.
- Szalma, J. L., J. S. Warm, G. Matthews, W. N. Dember, E. M. Weiler, A. Meier, and F. T. Eggemeier. 2004. Effects of sensory modality and task duration on performance, workload, and stress in sustained attention. *Hum Factors* 46:219–33.
- Thayer, R. E. 1989. *The Biopsychology of Mood and Arousal*. New York: Oxford University Press.
- Thropp, J. E., J. L. Szalma, and P. A. Hancock. 2004. Performance operating characteristics for spatial and temporal discriminations: Common or separate capacities? *Proc Hum Factors Ergon Soc* 48:1880–4.
- van Reekum, C., and K. R. Scherer. 1997. Levels of processing in emotion-antecedent appraisal. In *Cognitive Science Perspectives on Personality and Emotion*, ed. G. Matthews, 259–300. Amsterdam: Elsevier.
- Venkatesh, V. 2000. Determinants of perceived ease of use: Integrating control, intrinsic motivation, and emotion into the technology acceptance model. *Inf Syst Res* 11:342–65.
- Vicente, K. J., and J. Rasmussen. 1992. Ecological interface design: Theoretical foundations. *IEEE Trans Syst Man Cybern* 22:589–606.
- Wachtel, P. L. 1967. Conceptions of broad and narrow attention. *Psychol Bull* 68:417–29.
- Warm, J. S., W. N. Dember, and P. A. Hancock. 1996. Vigilance and workload in automated systems. In *Automation and Human Performance: Theory and Applications*, ed. R. Parasuraman and M. Mouloua, 183–200. Hillsdale, NJ: Erlbaum.
- Warm, J. S., R. Parasuraman, and G. Matthews. 2008. Vigilance requires hard mental work and is stressful. *Hum Factors* 50:433–41.
- Watson, D., and A. Tellegen. 1985. Toward a consensual structure of mood. *Psychological Bulletin* 98:219–235.
- Wickens, C. D. 1980. The structure of attentional resources. In *Attention and Performance VIII*, ed. R. Nickerson, 239–57. Hillsdale, NJ: Erlbaum.
- Wickens, C. D. 1984. Processing resources in attention. In *Varieties of Attention*, ed. R. Parasuraman and D. R. Davies, 63–102. New York: Academic Press.
- Wickens, C. D. 1996. Designing for stress. In *Stress and Human Performance*, ed. J. E. Driskell and E. Salas, 279–95. Mahwah, NJ: Erlbaum.
- Yeh, Y., and C. D. Wickens. 1988. Dissociation of performance and subjective measures of workload. *Hum Factors* 30:111–20.
- Yerkes, R. 1918. Psychology in relation to the war. *Psychol Rev* 25:85–115.
- Young, M. S., and N. A. Stanton. 2002. Malleable attentional resources theory: A new explanation for the effects of mental underload on performance. *Hum Factors* 44:365–75.
- Zhang, T., D. B. Kaber, B. Zhu, M. Swangnetr, P. Mosaly, and L. Hodge. 2010. Service robot feature design effects on user perceptions and emotional responses. *Intell Serv Rob* 3:73–88.
- Zimmerman, B. J. 2000. Attaining self-regulation: A social cognitive perspective. In *Handbook of Self-Regulation*, ed. M. Boekaerts, P. R. Pintrich, and M. Zeidner, 13–39. San Diego: Academic Press.

5 Choices and Decisions of Computer Users

Anthony Jameson

CONTENTS

5.1	Introduction	77
5.1.1	Concepts and Goals	77
5.1.2	Relationships to Other Human–Computer Interaction-Related Research.....	78
5.1.2.1	Interaction Design Guidelines and Principles; Help and Training	78
5.1.2.2	Recommender Systems.....	79
5.1.2.3	Persuasive Technology	79
5.1.3	Preview of Aspects of Preferential Choice and Decision Making	79
5.2	General Preferential Choice Problems	80
5.2.1	Decision about Whether to Use a Given System	80
5.2.2	Choice of a Method.....	81
5.2.3	Configuration Decision.....	81
5.3	Focusing on Goals and Values.....	81
5.3.1	What Constitutes a Good Choice or Decision?	81
5.3.2	Current Goals and Values	82
5.4	Situation Assessment and Option Identification	82
5.5	Anticipation of Consequences	83
5.5.1	Anticipating Experience	83
5.5.2	Anticipating the Consequences of Configuration Choices	83
5.6	Intertemporal Choice	84
5.6.1	Time Discounting	84
5.6.2	Choice Bracketing.....	85
5.7	Reuse of Previous Choices.....	85
5.7.1	Recognition-Primed Decision Making.....	85
5.7.2	Coherent Arbitrariness	86
5.7.3	Choices Based on Habit	86
5.7.4	The Role of Skill Acquisition	86
5.7.5	Example from Research on Method Selection	87
5.7.6	Concluding Remarks on the Reuse of Past Choices	87
5.8	Social Influence	87
5.9	Learning from Experience.....	88
5.10	Concluding Remarks	89
	Acknowledgments.....	89
	References.....	90

5.1 INTRODUCTION

5.1.1 CONCEPTS AND GOALS

Computer users are constantly making small choices and larger decisions about how to use their computing technology, such as the following:

- Which of the available photo management apps shall I use on my smartphone?

- Shall I dictate this e-mail message using speech recognition or tap in the text with a stylus?
- How should I configure my privacy settings?

This chapter focuses on cases, like these, in which a user can choose among two or more *options*, none of which is correct or incorrect but one of which can be *preferred* to the others. The term *preferential choice* will be used to distinguish this situation from *nonpreferential* choices that concern the correct way to operate a system, such as “Which of

these unfamiliar icons do I have to click on to send off my e-mail message?”

We will use the terms *choice* and *decision*, together and in alternation, to do justice to the variety of forms that the processes in question can take. *Decision* suggests a thorough, effortful process, whereas *choice* suggests a quick selection that may be based, for example, on habit. Both types of process occur in computer users, often with regard to the same set of options.

The following are the goals of this chapter:

1. Bring preferential choices and decisions of computer users into the foreground as a topic in human–computer interaction (HCI).
2. Provide access to the relevant psychological and HCI literature by summarizing key concepts and results and listing references.
3. Provide a framework for thinking about how to help computer users make better preferential choices and decisions.

5.1.2 RELATIONSHIPS TO OTHER HUMAN–COMPUTER INTERACTION-RELATED RESEARCH

Figure 5.1 visualizes the relationships between these goals and the goals of three other broad types of research that fall within or overlap with the HCI field.




FIGURE 5.1 Visualization of the relationships between the focus of this chapter and three human–computer interaction-related areas of research.

5.1.2.1 Interaction Design Guidelines and Principles; Help and Training

Much of what is known about how to design interactive systems and their associated help and training material can be seen as concerning ways of helping users to make the right choices: to click on the right icon or web link, select the correct command from a menu, or identify the part of the system that will provide the needed functionality. Interaction designers have become skilled at helping users to make these choices well, for example, by designing effective visual displays, making the user’s options clearly identifiable and understandable, providing informative feedback on the user’s actions, and making the actions reversible in case they do not yield a satisfactory result (see, e.g., Johnson 2010, for a collection of well-known sets of user-interface design guidelines). Similarly, those who develop online help and training programs have worked out a rich set of best practices for instructing and advising users about the choices that they need to make. Most of the content of help and training concerns the general question of how to operate the system in question, but some of it explicitly addresses preferential choices, such as when to use each of two available methods for accomplishing a particular goal or what type of configuration is best under what circumstances (e.g., “This setting is recommended if you often work off-line”).

Guidelines and design principles are rarely tailored explicitly to supporting preferential choices and decisions, and the related research hardly ever refers to the psychological literature on these topics that is covered in this chapter.

5.1.2.2 Recommender Systems

A focus on preferential choice and decisions is found, by contrast, in research on recommender systems (see, e.g., Jannach et al. 2011; Ricci et al. 2010) which aim to support and influence users' choices concerning products to buy, documents to read, and a variety of other types of items. As Figure 5.1 shows, recommender systems almost always support decisions that are not about the use of computing technology as such. The work in this area tends to be based to some extent on knowledge about psychological processes involved in preferential choice, but the main focus of attention is on accurately predicting what items will satisfy a user, rather than on understanding and influencing the user's decision-making processes.

5.1.2.3 Persuasive Technology

Yet another line of research (see, e.g., Fogg 2003; Fogg, Cueller, and Danielson 2008) differs from the previous paradigm mainly in its emphasis on motivating and persuading people to do some particular thing (e.g., save energy), which either that person or someone else has decided is best for the person in question. This line of research has yielded a wealth of ideas about how computing technology can be deployed to influence people's beliefs and behaviors. But only a few of the choices and behaviors targeted for persuasion (e.g., none of the 12 "domains for persuasive technology" listed in Table 7.1 of Fogg, Cueller, and Danielson 2008) concern computer use as such.

As Figure 5.1 indicates, this chapter will not go into much depth on the question of how to support and influence preferential choices concerning computer use. Instead, by foregrounding this class of choices and by providing an introduction to the large areas of relevant psychological literature, it aims to encourage and support increased attention to this topic.* Systematic efforts to support choices and decisions of this type should be able to benefit greatly from appropriately adapted knowledge transferred from the other three areas of research, notwithstanding the various differences visualized in Figure 5.1.

5.1.3 PREVIEW OF ASPECTS OF PREFERENTIAL CHOICE AND DECISION MAKING

Figure 5.1 reflects the fact that psychological research about how people make preferential choices and decisions has received limited attention in HCI so far.† One reason may be the fact that there is no single relevant theory in psychology that could be straightforwardly adapted to the needs of the HCI

field. Although dozens of books and hundreds of articles from relevant psychological research exist, they come from several research traditions that only partly overlap and refer to each other. The discussion in this chapter will draw from these areas: judgment and decision making (see, e.g., Hastie and Dawes 2010; Koehler and Harvey 2004; Lichtenstein and Slovic 2006; Schneider and Shanteau 2003; Newell, Lagnado, and Shanks 2007; Weber and Johnson 2009); naturalistic decision making (Klein 1998), the reasoned action approach (Fishbein and Ajzen 2010), research on habitual behavior (Wood and Neal 2007), behavioral economics (Ariely 2008; Iyengar 2010; Thaler and Sunstein 2008), and research on self-control (Rachlin 2000) and on compliance tactics (Cialdini 2007).

As a way of providing a reasonably coherent overview despite the differences among these research traditions and their terminologies, Table 5.1 lists the aspects of choice and decision processes that will be covered in turn in this chapter, formulating each one in terms of one or more "questions" that a computer user might conceivably "ask" him- or herself while considering a choice or decision. Although in some cases such questions may be consciously asked and addressed by a computer user, the processing represented in the table by a question often occurs without any verbal formulation or conscious deliberation—whatever particular definition of the elusive concept of *consciousness* one may prefer to use (see, e.g., Wilson 2002).

With any given choice or decision for a particular person, in general only some subset of these considerations will be relevant, and the table is not intended to convey a particular temporal order of processing: Because of the variety of forms that

TABLE 5.1
Preview of the Aspects of Preferential Choice and Decision Making Discussed in This Chapter

Topic	Questions That a Decision Maker May Consider
Focusing on goals and values	What is a good decision-making process for this situation? What are my relevant goals and values?
Situation assessment and option identification	What's going on in this situation? What are my options?
Anticipation of consequences	What would the consequences be if I chose this option? How desirable would they be?
Intertemporal choice	How should I value consequences that will not occur until sometime in the future? How should I deal with a sequence of repetitions of basically the same choice?
Reuse of previous choices	What did I choose the last time I had a choice like this?
Social influence	What do other people choose in this situation? What do they want or expect me to choose?
Learning from experience	What can I learn from the results of the choice that I have made?

* A first step toward a systematic approach to supporting preferential choice on the basis of the conceptual framework of this chapter is offered by Jameson et al. (2011).

† Two thorough book-length syntheses of cognitive psychology research for HCI (Gardiner and Christie 1987; Johnson 2010) include hardly any references to the sort of psychology literature cited in this chapter.

preferential choices and decisions can take, it would not be realistic to try to formulate a causal model or a process model, for example, in the form of a flowchart, though models of this sort are often found useful for particular types of choice or decision-making situation (see, e.g., Wickens and Hollands 2000, Chapter 7; Fishbein and Ajzen 2010; Klein 1998, Chapter 3).

5.2 GENERAL PREFERENTIAL CHOICE PROBLEMS

Although opportunities to make preferential choices and decisions crop up constantly with just about every type of interactive system, there are three generic classes of choice that are worth distinguishing, because of their frequency of occurrence and because they have attracted a fair amount of attention in HCI research. Table 5.2 introduces them to facilitate reference to them at various points later in the chapter.

5.2.1 DECISION ABOUT WHETHER TO USE A GIVEN SYSTEM

One type of decision that a person can make with regard to computer use is that of whether to use a given system at all. The most extensive line of research that has looked into this question is research on *technology acceptance*. A good entry point to this literature is the influential article by Venkatesh et al. (2003), which presented the Unified Theory of Acceptance and Use of Technology (UTAUT), a model that integrates eight previously developed models, including the especially widely studied *technology acceptance model* (see, e.g., Venkatesh and Davis 2000). These models in turn drew their inspiration from more general theories from social psychology and sociology, such as the precursors of the recently formulated *reasoned action approach* of Fishbein and Ajzen (2010).

TABLE 5.2
Three General Types of Preferential Choice That Have Been Studied in Human–Computer Interaction

Generic Choice Problem	Selected Research Issues
Decision about whether to use a given system	What variables influence people’s decisions about whether to use a given system if it is made available to them (usually: within an organization)? What are the causal relationships among these variables? How can these variables be measured?
Choice of a method from a set of alternative methods	When more than one method is available for a particular subtask, how do users decide which one to use? Why do even experienced users sometimes persist in using inefficient methods?
Configuration decision	How do people decide whether and when to configure an application? What difficulties do they encounter when making configuration choices?

Table 5.3 gives an impression of the basic nature of the models in this area by depicting the four main variables in the UTAUT model that influence intention to use a given system and actual use of the system, along with examples of questionnaire items typical of those used to measure these variables. The model also includes claims about several variables that moderate the influence of these main variables: *gender*, *age*, *experience*, and *voluntariness of use*.

Although some of these questions are reminiscent of questions from usability scales such as System Usability Scale (SUS) (Brooke 1996), the overall goal of the model and the associated measuring instruments are not to assess usability but rather to predict whether potential users (typically, employees in a given company) will actually use a given system (e.g., a new videoconferencing system) if it is made available to them. Note that most of the questions related to the variables *social influence* and *facilitating conditions* concern considerations other than usability.

Researchers and practitioners in the HCI field usually want to go beyond *predicting* whether people in a given target group will use a given (type of) system, to attempt to improve the system (and/or related resources) to increase the likelihood that the system will be used and the success of its use. Still, the large amount of information collected in the technology acceptance area about variables related to choices about system use and about ways of measuring these variables can help to stimulate and structure thinking about this class of choices. Researchers in this area regularly introduce new variables and new perspectives that shed light on different aspects of acceptance decisions (see, e.g., Bagozzi 2007; Loraas and Diaz 2009).

TABLE 5.3
The Four Main Variables in the UTAUT Model and Typical Questionnaire Items Used to Measure Them

<i>Performance expectancy</i>
Using the system in my job would enable me to accomplish tasks more quickly.
Using the system would improve my job performance.
Using the system would make it easier to do my job.
<i>Effort expectancy</i>
Learning to operate the system would be easy for me.
My interaction with the system would be clear and understandable.
I would find the system to be flexible to interact with.
<i>Social influence</i>
People who influence my behavior think that I should use the system.
People who are important to me think that I should use the system.
<i>Facilitating conditions</i>
I have control over using the system.
I have the resources necessary to use the system.
I have the knowledge necessary to use the system.
The system is not compatible with other systems I use.

Source: Based on parts of Figure 3 and Tables 9–12 of Venkatesh, V., M. G. Morris, G. B. Davis, and F. D. Davis. 2003. *MIS Quart* 27(346): 425–78.

5.2.2 CHOICE OF A METHOD

In all but the simplest interactive systems, there is often more than one method available for achieving a given goal. Whenever the user can choose freely between two or more methods, the choice is preferential. Card, Moran, and Newell (1983) introduced in their Goals, Operators, Methods, and Selection Rules (GOMS) model (described most completely in Card, Moran, and Newell 1983; see also Kieras 2008) a notation for such cases: the two or more available methods are described as part of the model for a given task, and it is assumed that each user has learned a selection rule for making the choice (e.g., “Use the mouse instead of the cursor keys if the target is more than a couple of inches away on the screen”); this assumption is plausible given that the GOMS model assumes that users have considerable experience with the system and the tasks in question.

In the intervening years, some research has looked at the ways in which users learn selection rules on the basis of experience with the methods in question (see, e.g., Gray and Boehm-Davis 2000) and at the considerations that users take into account when choosing among methods (see, e.g., Young and MacLean 1988; Jameson and Klöckner 2005), whereas other researchers have investigated situations in which users systematically fail to use suitable methods that are available to them (Carroll and Rosson 1987; Bhavnani and John 2000; Bhavnani, Peck, and Reif 2008; Charman and Howes 2003).

5.2.3 CONFIGURATION DECISION

A usually more complicated type of choice that users can make concerns whether, when, and how to configure an application to suit their own tastes and needs. Over the years, researchers have repeatedly found this type of problem to be challenging for most users (see, e.g., Mackay 1991; McGrenere, Baecker, and Booth 2007), and it has attracted increased attention in recent years because of the practically important problem of configuring privacy settings in social network platforms (see, e.g., Iachello and Hong 2007).

5.3 FOCUSING ON GOALS AND VALUES

The first of the general considerations listed in Table 5.1 concerns the basic values that a chooser will be guided by when making a choice. Although computer users often do not think explicitly about these values, interaction designers ought to be aware of them when considering how to support good choices; and calling these issues to the user’s attention may be an effective tactic.

5.3.1 WHAT CONSTITUTES A GOOD CHOICE OR DECISION?

The most fundamental question is that of what constitutes a good choice in the first place. Before considering what choosers think about this issue, we should notice a shift in the thinking of scientists who have studied decision making. Traditional notions of what constitutes a good decision are that a decider should (1) apply a decision procedure

that is normatively justifiable (e.g., consistent with the laws and principles of logic, probability, and expected utility) and (2) choose the action that will maximize desirable (and minimize undesirable) outcomes under idealized conditions (see, e.g., Gigerenzer and Todd 1999, Chapter 1; Gigerenzer 2007, Chapter 5). More recently, researchers have become impressed by the extent to which animals and humans can function quite effectively by using decision procedures that are justifiable only in the sense that they work well in the environment in which they are applied and make good use of the decider’s limited time and cognitive resources. For example, a web searcher’s strategy of clicking on the first link on the search result page that looks reasonably relevant would be hard to justify in terms of a normatively optimal general strategy; but if the user’s previous experience with the search engine in question has shown that the first reasonably relevant-looking link is almost always the best one, this strategy can be considered *ecologically rational* for that search engine. The same point can apply to the decision rule of always buying your smartphone applications from your favorite vendor or always accepting the default configuration when installing new software. In cases where the choices of a computer user make sense only given particular assumptions about the structure of the environment, the best way to help the user make good choices may be to ensure that the environment fulfills these assumptions.

Researchers have also investigated the question of what constitutes a good decision process from the point of view of the decision maker (see, e.g., Bettman, Luce, and Payne 2006; Hastie 2001; Yates, Veinott, and Patalano 2003). Although specific answers to this question vary, the following statements are widely accepted:

1. Choosers want their decision to yield a good outcome.

This point is not as straightforward as it may seem, because what counts as a good outcome depends in turn on a variety of factors, as we will see.

2. Choosers do not want to invest time and effort in the decision-making process itself that is out of proportion to the benefits of doing so.

For example, when installing a new application, a user who is asked which specific components should be installed may choose the option “Everything” simply to save the time of deciding about the individual components, since the possible benefits of choosing any other option (e.g., saving a few megabytes of hard disk space) do not seem to justify the investment of even a few seconds of decision time.

3. Choosers prefer to avoid unpleasant thoughts.

Some ways of thinking about a decision can involve distressing thoughts, as when a driver faces a choice between (1) ignoring an incoming text message from his boss and (2) driving less safely for a while in order to respond to the message. A user may be motivated to think about the decision in a way that avoids such thoughts (e.g., by convincing

himself that he can respond to the boss’s message without taking the slightest risk).

4. Choosers often want to be able to justify the decision that they have made to other persons—or to themselves.

Justifiability is often simply a necessary condition for being able to implement a decision (cf. Lerner and Tetlock 2003). For example, even if a business person would really like to buy an iPhone for professional use, they may prefer a Blackberry instead because they think that this choice is more likely to be approved by their company’s purchasing department. But even just the desire to convince another person or oneself that a decision was sound can cause people to look for justifiable decisions (see, e.g., Shafir, Simonson, and Tversky 2006).

Consequently, one way of supporting preferential choice is to make it easy for the user to come up with a satisfying justification of whatever option is best for him, for example, by supplying a justification explicitly (as is done by many recommender systems) (see Tintarev and Masthoff 2010) or by structuring the situation in such a way that a justification is easy to derive.

5.3.2 CURRENT GOALS AND VALUES

One characteristic of preferential choice is its dependence on the particular goals that the chooser is currently focusing on (see, e.g., Schneider and Barnes 2003). To a certain extent, this dependence is obviously necessary and appropriate: Your choice of an application to prepare a text document with should depend on whether you want it to be beautifully formatted or whether you just want to get it finished as quickly as possible. But the dependence on current goals can also lead to some curious phenomena: Both anecdotal evidence and some research (e.g., Iachello and Hong 2007, Section 3.3.2; Mackay 1991) concerning configuration decisions tell us that users often accept the default configuration of a system until some negative event (e.g., a privacy violation or a need to repeat a given tedious operation multiple times) prompts them to change the configuration. A normatively more rational way of deciding when and what to configure would involve something like estimating the total (discounted) benefit of the improved configuration over an extended period of system use. In contrast, reactive configuration can be seen as a response to the goal of preventing the specific negative thing that just happened from ever happening again. Whether this configuration action is really a good idea in the long run will depend on how well the short-term goal happens to coincide with the user’s larger pattern of goals and use situations. Mackay (1991) and Iachello and Hong (2007) offer perceptive discussions of strategies for dealing with this type of discrepancy.

Keeney (1992) discusses in great depth the importance of ensuring that decisions depend on the decision maker’s true values rather than on temporarily salient considerations

such as those that are suggested by the set of options that are immediately available. Although interaction designers rarely, if ever, have an opportunity to support their users with in-depth decision analysis, calling the user’s attention to important goals and values on a much smaller scale does represent a promising way of supporting preferential choice. Two experiments by Mandel and Johnson (2006) demonstrate clearly how a goal or value (e.g., “safety” or “economy” for a prospective car buyer) can be activated by a change in interface design (e.g., the colored background of the web pages of an e-commerce site), mostly without awareness on the part of the user.

5.4 SITUATION ASSESSMENT AND OPTION IDENTIFICATION

To be able to make a choice or decision, the chooser must normally in some sense be aware of the fact that a choice is available—though in extreme cases the awareness can be minimal, as when the choice is made out of habit or when it involves accepting the status quo or default option by doing nothing.

In experimental laboratory studies, the way in which the chooser perceives or “frames” the choice problem is largely under the control of the experimenter. Some well-known and striking results concern the effects on choice of the way in which the problem is framed. For example, people tend to be influenced strongly by whether options are described in terms of people being “saved” versus people “dying,” even when the situations described in these terms are objectively identical. An important part of one of the dominant theories of judgment and decision making, *prospect theory* (originally presented by Kahneman and Tversky 1979), concerns the process of “editing” the initial representation of a choice problem to arrive at the chooser’s own representation; but choosers often stick with the initial representation.

Like laboratory experimenters, interaction designers often have control over the way in which a choice is presented to the user. For example, users who purchase a software product are often offered an option like “Check this box to receive news about updates and special offers,” which a user may mentally edit into a representation like “Check this box to get even more spam.”

When decision making occurs outside the laboratory, the presentation of the choice problem is often less clear-cut; understanding the situation and identifying the available options can be a complex process (often called *situation assessment*) that calls for considerable expertise. This process has been extensively studied within the research paradigm of *naturalistic decision making* (see, e.g., Klein 1998; Klein 2008; Maule 2010). This type of decision making is typified by the situation of a fire brigade arriving at the scene of a burning building: The problem situation is changing rapidly over time, even as the decision makers contemplate how to deal with the fire; there is considerable stress because of the high stakes and because of environmental factors such as noise and heat; and on the positive side, the decision makers typically possess considerable experience in dealing with

such situations, which makes it unnecessary for them to analyze the problem from first principles. Some key results of this research will be summarized below in Section 5.7.1. For now, the main point is that recognizing the need for a choice and identifying or generating one or more options is sometimes the most important and challenging aspect of a decision problem.

An implication for interaction design is that we should look out for situations in which recognizing and interpreting a decision situation may be unnecessarily or unduly challenging for at least some users. For example, a sophisticated user who installs a new web browser is likely to recognize the need to choose security and privacy settings that are well adapted to the context in which the browser will be used; a less sophisticated user is likely to accept the default settings, perhaps without even being aware that a choice exists.

In fact, the widespread tendency of people to overlook or ignore choice opportunities and accept the default represents a major way in which *choice architects* (to use the suggestive term of Thaler and Sunstein 2008), including interaction designers, can influence choices. Widely discussed controversies concerning computer use include the bundling of software with the Windows operating system (which offers new users a convenient default option for many application choices that they would otherwise have to make) and the default privacy options for social network platforms like Facebook. Outside of the arena of computer use, one of the primary and most successful tactics of interventions based on behavioral economics (such as the *libertarian paternalism* of Thaler and Sunstein 2008) is to provide a default option which is thought to be in the best interest of the people making the choice in question or of society as a whole (e.g., laws that state that every person can be viewed as an organ donor unless they have specified otherwise; see Johnson and Goldstein 2006).

5.5 ANTICIPATION OF CONSEQUENCES

The most dominant traditional view of decision making is a *consequentialist* one: that of a person who contemplates the (perhaps uncertain) consequences of choosing each of the available options and bases the decision on an evaluation of those consequences. As Table 5.1 indicates, there are other considerations that can affect a decision, and in fact, choosers sometimes do not contemplate consequences at all.

Still, computer users do sometimes anticipate the possible consequences of their choices, and one question that arises is that of what sorts of consequence they anticipate. If computer users were concerned only about traditional usability criteria, they might make their decisions solely on the basis of consideration of consequences like those covered by UTAUT's *performance expectancy* and *effort expectancy* variables (Table 5.3). The growing interest in recent years in a broader view of user experience (see, e.g., Law et al. 2009; Kuniavsky 2010) can be viewed as an awareness of a wider range of types of consequence that can influence users' evaluations of systems and possible actions.

5.5.1 ANTICIPATING EXPERIENCE

But how accurately can computer users anticipate the consequences of options? Even just anticipating the enjoyableness of an experience that has been described to you (e.g., using an allegedly delightful photo management app on a smartphone) is not as straightforward as it would intuitively seem. Trying the experience out briefly (e.g., with a demo version of the app) is not always a reliable test, partly because of people's tendency to adapt their tastes and expectations on the basis of new experience (see, e.g., Wilson 2002, Chapter 7). And if a user's initial expectation is (erroneously) that an experience will not be positive, he or she may refrain from trying it out in the first place.

A straightforward effort of designers to support the anticipation of the experience of performing an action is found in promises such as "Filling in our customer satisfaction questionnaire will take just 2 minutes of your time" or "Configuring the application is quick and easy." But this method presupposes that the user is likely to believe claims like these. An alternative approach is to consider nonverbal ways of previewing the consequences of an action. This general strategy has been explored extensively in the area of persuasive technology (see, e.g., Fogg 2003, Chapter 4), as with the "Baby Think it Over" infant simulator, which helps teen-aged girls anticipate realistically what it is like to take care of a baby. Some further work will probably be required before this strategy can be applied widely to (1) decisions concerning computer use and (2) decisions where it is not a priori clear which option is best for the chooser—that is, where the chooser must really *choose*, as opposed to being persuaded (cf. Figure 5.1).

5.5.2 ANTICIPATING THE CONSEQUENCES OF CONFIGURATION CHOICES

One challenge for users in connection with the configuration of applications (Mackay 1991; Iachello and Hong 2007) is that the consequences of configuration actions tend to be hard to anticipate. First, there is the question of how time-consuming, tedious, and risky the configuration actions themselves will be. Then there is the fact that the consequences of a configuration decision are often not immediately visible; they consist in changes to the computing environment that will have consequences in the future which will in turn depend on actions of the user and other configuration settings.

Gabrielli and Jameson (2009), applying an adapted heuristic walkthrough to parts of four widely used applications, found that about three-quarters of the formulations used to describe configuration options (e.g., "Accept cookies from third parties") did not appear to convey to a typical user a clear idea of the meaning of an option, the consequences of choosing it, or the overall desirability of choosing it. The proportion of problematic cases diminished to about one-half if the help texts explaining the options were taken into account.