

OXFORD LIBRARY OF PSYCHOLOGY

EDITED BY

**JOHN D.
LEE &**

**ALEX
KIRLIK**



≡ The Oxford Handbook *of*
**COGNITIVE
ENGINEERING**



OXFORD LIBRARY OF PSYCHOLOGY

Editor in Chief PETER E. NATHAN

The Oxford Handbook of Cognitive Engineering

Edited by

John D. Lee

Alex Kirlik

Consulting Editors

Marvin J. Dainoff

Peter A. Hancock

Neville Moray

Donald A. Norman

Richard W. Pew

Eduardo Salas

Thomas B. Sheridan

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide.

Oxford New York
Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in
Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trademark of Oxford University Press in the UK and certain other countries.

Published in the United States of America by
Oxford University Press
198 Madison Avenue, New York, NY 10016

© Oxford University Press 2013

Oxford is a registered trademark of Oxford University Press

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, without the prior permission in writing of Oxford University Press, or as expressly permitted by law, by license, or under terms agreed with the appropriate reproduction rights organization. Inquiries concerning reproduction outside the scope of the above should be sent to the Rights Department, Oxford University Press, at the address above.

You must not circulate this work in any other form
and you must impose this same condition on any acquirer.

Library of Congress Cataloging-in-Publication Data
The Oxford handbook of cognitive engineering / edited by John D. Lee, Alex Kirlik.
pages cm. — (Oxford library of psychology)
ISBN 978-0-19-975718-3 (hardback)
1. User-centered system design. 2. Human-computer interaction. 3. Human engineering.
I. Lee, John D., 1965- editor of compilation. II. Kirlik, Alex, editor of compilation.
QA76.9.H85O94 2013
004.01'9—dc23
2012026107

9 8 7 6 5 4 3 2 1
Printed in the United States of America
on acid-free paper

SHORT CONTENTS

Oxford Library of Psychology vii

About the Editors ix

Contributors xi

Table of Contents xv

Chapters 1–622

Index 623

This page intentionally left blank

The *Oxford Library of Psychology*, a landmark series of handbooks, is published by Oxford University Press, one of the world's oldest and most highly respected publishers, with a tradition of publishing significant books in psychology. The ambitious goal of the *Oxford Library of Psychology* is nothing less than to span a vibrant, wide-ranging field and, in so doing, to fill a clear market need.

Encompassing a comprehensive set of handbooks, organized hierarchically, the *Library* incorporates volumes at different levels, each designed to meet a distinct need. At one level are a set of handbooks designed broadly to survey the major subfields of psychology; at another are numerous handbooks that cover important current focal research and scholarly areas of psychology in depth and detail. Planned as a reflection of the dynamism of psychology, the *Library* will grow and expand as psychology itself develops, thereby highlighting significant new research that will impact the field. Adding to its accessibility and ease of use, the *Library* will be published in print and, later on, electronically.

The *Library* surveys psychology's principal subfields with a set of handbooks that capture the current status and future prospects of those major subdisciplines. This initial set includes handbooks of social and personality psychology, clinical psychology, counseling psychology, school psychology, educational psychology, industrial and organizational psychology, cognitive psychology, cognitive neuroscience, methods and measurements, history, neuropsychology, personality assessment, developmental psychology, and more. Each handbook undertakes to review one of psychology's major subdisciplines with breadth, comprehensiveness, and exemplary scholarship. In addition to these broadly conceived volumes, the *Library* also includes a large number of handbooks designed to explore in depth more specialized areas of scholarship and research, such as stress, health and coping, anxiety and related disorders, cognitive development, or child and adolescent assessment. In contrast to the broad coverage of the subfield handbooks, each of these latter volumes focuses on an especially productive, more highly focused line of scholarship and research. Whether at the broadest or most specific level, however, all of the *Library* handbooks offer synthetic coverage that reviews and evaluates the relevant past and present research and anticipates research in the future. Each handbook in the *Library* includes introductory and concluding chapters written by its editor to provide a roadmap to the handbook's table of contents and to offer informed anticipations of significant future developments in that field.

An undertaking of this scope calls for handbook editors and chapter authors who are established scholars in the areas about which they write. Many of the nation's and the world's most productive and best-respected psychologists have agreed to edit *Library* handbooks or write authoritative chapters in their areas of expertise.

For whom has the *Oxford Library of Psychology* been written? Because of its breadth, depth, and accessibility, the *Library* serves a diverse audience, including graduate students in psychology and their faculty mentors, scholars, researchers, and practitioners in psychology and related fields. Each will find in the *Library* the information they seek on the subfield or focal area of psychology in which they work or are interested.

Befitting its commitment to accessibility, each handbook includes a comprehensive index, as well as extensive references to help guide research. And because the *Library* was

designed from its inception as an online as well as a print resource, its structure and contents will be readily and rationally searchable online. Further, once the *Library* is released online, the handbooks will be regularly and thoroughly updated.

In summary, the *Oxford Library of Psychology* will grow organically to provide a thoroughly informed perspective on the field of psychology, one that reflects both psychology's dynamism and its increasing interdisciplinarity. Once published electronically, the *Library* is also destined to become a uniquely valuable interactive tool, with extended search and browsing capabilities. As you begin to consult this handbook, we sincerely hope you will share our enthusiasm for the more than 500-year tradition of Oxford University Press for excellence, innovation, and quality, as exemplified by the *Oxford Library of Psychology*.

Peter E. Nathan
Editor-in-Chief
Oxford Library of Psychology

ABOUT THE EDITORS

John D. Lee

John D. Lee is the Emerson Electric professor in the Department of Industrial and Systems Engineering at the University of Wisconsin, Madison, and director of the Cognitive Systems Laboratory. Previously he was a professor at the University of Iowa and director of human factors research at the National Advanced Driving Simulator. He is a coauthor of the textbook *An Introduction to Human Factors Engineering* and is the author or coauthor of over 170 articles. He recently helped edit *The Oxford Handbook of Cognitive Engineering*. He has served on the National Academy of Sciences Committee on Human System Integration, the Committee on Electronic Vehicle Controls and Unintended Acceleration, and several other committees. His research focuses on the safety and acceptance of complex human-machine systems by considering how technology mediates attention. He now serves on the editorial board of *Cognitive Engineering and Decision Making*; *Cognition, Technology and Work*; and the *International Journal of Human Factors Modeling and Simulation*, and is the associate editor for the journals *Human Factors* and *IEEE-Systems, Man, and Cybernetics*. He received the Ely Award for best paper in the journal *Human Factors* (2002) and the best paper award for the journal *Ergonomics* (2005). Specific research interests include trust in technology, advanced driver assistance systems, and human interaction with mixed-initiative systems.

Alex Kirlik

Alex Kirlik is professor in the Departments of Computer Science, Psychology, and Systems Engineering and a member of the Beckman Institute at the University of Illinois at Urbana-Champaign. He previously served as Coordinator of Cognitive Science Academic Programs at Georgia Institute of Technology and as Head of Human Factors at Illinois. He is the editor of *Adaptive Perspectives on Human-Technology Interaction* (Oxford, 2006), co-editor of *Attention: From Theory to Practice* (Oxford, 2008), and editor of *Human-Tech* (Oxford, 2011). His engineering research focuses on the design of collaborative human-technology systems to leverage mathematical models and human expertise in cognitive tasks such as judgment, prediction, diagnosis, and decision making. He is a member of the editorial boards of the journals *Human Factors* and *Cognitive Engineering and Decision Making* and has served as editor of the Oxford Series in Human-Technology Interaction since its inception in 2006.

This page intentionally left blank

CONTRIBUTORS

Leonard Adelman

Department of Systems Engineering and
Operations Research
The Volgenau School of Engineering
George Mason University
Fairfax, VA

Arna Banerjee

Departments of Anesthesiology, Critical
Care and the Vanderbilt Center for
Perioperative Research in Quality
Vanderbilt University Medical Center
Nashville, TN

Gordon D. Baxter

Department of Computer Science
University of St. Andrews
St. Andrews, Scotland

Aaron S. Benjamin

Department of Psychology
University of Illinois at Urbana-Champaign
Urbana-Champaign, IL

Kevin B. Bennett

Department of Psychology
Wright State University
Dayton, OH

Ann M. Bisantz

Industrial and Systems Engineering
University at Buffalo
Buffalo, NY

Walter R. Boot

Department of Psychology
Florida State University
Tallahassee, FL

Catherine M. Burns

Systems Design Engineering
University of Waterloo
Waterloo, Ontario, Canada

Michael D. Byrne

Departments of Psychology and Computer
Science
Rice University
Houston, TX

Pascale Carayon

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

Anna T. Cianciolo

Southern Illinois University School of
Medicine
Command Performance Research, Inc.
Champaign, IL

Richard I. Cook

Department of Patient Safety School of
Technology and Health
Royal Institute of Technology
Stockholm, Sweden

Nancy J. Cooke

Department of Biomedical Informatics
Arizona State University
Scottsdale, AZ

Beth W. Crandall

Cognitive Solutions Division
Applied Research Associates, Inc.
Fairborn, OH

Mary L. Cummings

Department of Aeronautics and
Astronautics
Massachusetts Institute of Technology
Cambridge, MA

Asaf Degani

General Motors R&D
Advanced Technical Center
Israel

Sidney Dekker

Key Centre for Ethics, Law, Justice and
Governance
Griffith University
Southport, Queensland, Australia

Birsen Donmez

Department of Mechanical and Industrial
Engineering
University of Toronto
Toronto, Ontario, Canada

L. Jane Easdown

Departments of Anesthesiology, Critical
Care, and the Vanderbilt Center for
Perioperative Research in Quality
Vanderbilt University Medical Center
Nashville, TN

Mica R. Endsley

SA Technologies, Inc.

K. Anders Ericsson

Department of Psychology
Florida State University
Tallahassee, FL

Karen M. Evans

Rochester Institute of Technology
Rochester, NY

Stephen M. Fiore

Department of Philosophy
University of Central Florida
Orlando, FL

Ute M. Fischer

School of Literature, Communication, and
Culture
Georgia Institute of Technology
Atlanta, GA

Arthur D. Fisk

School of Psychology
Georgia Institute of Technology
Atlanta, GA

John M. Flach

Department of Psychology
Wright State University
Dayton, OH

Wai-Tat Fu

Human Factors
University of Illinois at Urbana-Champaign
Urbana-Champaign, IL

Gerd Gigerenzer

Center for Adaptive Behavior and
Cognition
Max Planck Institute for Human
Development
Berlin, Germany

Kevin A. Gluck

Cognitive Models and Agents Branch
Air Force Research Laboratory
Wright-Patterson Air Force Base, OH

Michael A. Goodrich

Computer Science Department
Brigham Young University
Provo, UT

Jamie C. Gorman

Department of Psychology
Texas Tech University
Lubbock, TX

Stephen J. Guastello

Department of Psychology
Marquette University
Milwaukee, WI

Glenn Gunzelmann

Cognitive Models and Agents Branch
Air Force Research Laboratory
Wright-Patterson Air Force Base, OH

Michael Heymann

Department of Computer Science
Israel Institute of Technology
Haifa, Israel

Jean-Michel Hoc

Centre National de la Recherche
Scientifique
Université Nantes Angers Le Mans
Institut de Recherche en
Communications et Cybernétique
de Nantes, France

Robert R. Hoffman

Institute for Human and Machine
Cognition
Pensacola, FL

Peter Hoonakker

Center for Quality and Productivity
Improvement
University of Wisconsin-Madison
Madison, WI

Daniel Hummerdal

MINES-ParisTech
Paris, France

Richard J. Jagacinski

Department of Psychology
The Ohio State University
Columbus, OH

David B. Kaber

Departments of Industrial & Systems
Engineering and Psychology
North Carolina State University
Raleigh, NC

Konstantinos V. Katsikopoulos

Max Planck Institute for Human
Development
Berlin, Germany

Jong W. Kim

University of Central Florida
Orlando, FL

Alex Kirlik

Departments of Computer Science,
Psychology, and Systems Engineering
University of Illinois at
Urbana-Champaign
Urbana, IL

Gary Klein

MacroCognition, LLC
Yellow Springs, OH

John D. Lee

Department of Industrial and Systems
Engineering
University of Wisconsin-Madison
Madison, WI

Yili Liu

Department of Industrial and Operations
Engineering
University of Michigan
Ann Arbor, MI

Jason S. McCarley

School of Psychology
Flinders University
Adelaide, South Australia, Australia

Joachim Meyer

Department of Industrial Engineering and
Management
Ben-Gurion University of the Negev
Beer-Sheva, Israel

Laura G. Militello

Applied Decision Science, LLC
Cincinnati, OH

Daniel G. Morrow

Beckman Institute for Advance Science and
Technology
University of Illinois at Urbana-Champaign
Urbana-Champaign, IL

Patricia Bockelman Morrow

Cognitive Sciences Laboratory
University of Central Florida
Orlando, FL

Kathleen L. Mosier

Department of Psychology
San Francisco State University
San Francisco, CA

Max Mulder

Department of Aerospace Engineering
Technical University of Delft
Delft, The Netherlands

Christopher P. Nemeth

Principal Scientist, Cognitive Solutions
Division
Applied Research Associates, Inc.
Albuquerque, NM

Marita A. O'Brien

Department of Psychology
University of Alabama in Huntsville
Huntsville, AL

Raja Parasuraman

Department of Psychology
George Mason University
Fairfax, VA

Peter Pirolli

Palo Alto Research Center
Palo Alto, CA

Amy R. Pritchett

School of Aerospace Engineering
Georgia Institute of Technology
Atlanta, GA

Frank E. Ritter

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

Wendy A. Rogers

School of Psychology
Georgia Institute of Technology
Atlanta, GA

Emilie M. Roth

Roth Cognitive Engineering
Brookline, MA

Dario D. Salvucci

Department of Computer Science
Drexel University
Philadelphia, PA

Nadine B. Sarter

Department of Industrial and
Operations Engineering
Center for Ergonomics
University of Michigan
Ann Arbor, MI

Michael Shafto

Intelligent Systems Division
NASA Ames Research Center
Moffett Field, CA

Sowmyalatha Srinivasmurthy

Department of Computer Science and
Engineering
The Pennsylvania State University
University Park, PA

Rene van Paassen

Aerospace Engineering
Technical University of Delft
Delft, The Netherlands

Matthew B. Weinger

Departments of Anesthesiology, Critical
Care, and the Vanderbilt Center for
Perioperative Research in Quality
Vanderbilt University Medical Center
Nashville, TN

Christopher D. Wickens

Human Factors
University of Illinois, Urbana-Champaign
Urbana-Champaign, IL

Alexander Wilhelmsson

Lund University
Lund, Sweden

CONTENTS

Part One • Cognitive Engineering: History and Foundations

Introduction to the Handbook 03

John D. Lee and Alex Kirlik

Part Two • Cognition in Engineered Systems

1. The Closed-Loop Dynamics of Cognitive Work 19
John M. Flach, Kevin B. Bennett, Richard J. Jagacinski, Max Mulder, and Rene van Paassen
2. Attention 36
Christopher D. Wickens
3. Multitasking 57
Dario D. Salvucci
4. Judgment and Prediction 68
Kathleen L. Mosier
5. Situation Awareness 88
Mica R. Endsley
6. Trust, Reliance, and Compliance 109
Joachim Meyer and John D. Lee
7. Learning and Retention 125
Frank E. Ritter, Gordon D. Baxter, Jong W. Kim, and Sowmyalatha Srinivasmurthy
8. Expertise 143
Walter R. Boot and K. Anders Ericsson
9. Neuroergonomics: Brain-Inspired Cognitive Engineering 159
Raja Parasuraman
10. Communication in Socio-Technical Systems 178
Daniel G. Morrow and Ute M. Fischer
11. Team Cognition: Coordination across Individuals *and* Machines 200
Patricia Bockelman Morrow and Stephen M. Fiore
12. Organizational Design and Cognitive Work 216
Pascale Carayon and Peter Hoonakker

Part Three • Cognitive Engineering Methods

13. Cognitive Task Analysis 229
Beth W. Crandall and Robert R. Hoffman
14. Cognitive Work Analysis 240
Emilie M. Roth and Ann M. Bisantz

15. Decision-Centered Design 261
Laura G. Militello and Gary Klein
16. Situation Awareness-Oriented Design 272
Mica R. Endsley
17. Cognitive Engineering to Support Successful Aging 286
Wendy A. Rogers, Marita A. O'Brien, and Arthur D. Fisk
18. Artifact Analysis as a Way to Understand Cognition 302
Christopher P. Nemeth and Richard I. Cook
19. Evaluation: Does the Cognitive Engineering Effort Do What It Was Envisioned to Do? 315
Leonard Adelman
20. Microworld Experimentation with Teams 327
Nancy J. Cooke and Jamie C. Gorman
21. Simulation to Assess Human Responses to Critical Events 336
L. Jane Easdown, Arna Banerjee, and Matthew B. Weinger
22. Simulation to Assess Safety in Complex Work Environments 352
Amy R. Pritchett
23. Metrics for Supervisory Control System Evaluation 367
Mary L. Cummings and Birsen Donmez
24. Multitasking and Multi-Robot Management 379
Michael A. Goodrich
25. Human-Machine Cooperation 395
Jean-Michel Hoc
26. Learning from Failure 404
Daniel Hummerdal, Alexander Wilhelmsson, and Sidney Dekker

Part Four • Cognitive Engineering Models

27. Computational Cognitive Modeling of Interactive Performance 415
Michael D. Byrne
28. Computational Process Modeling and Cognitive Stressors: Background and Prospects for Application in Cognitive Engineering 424
Kevin A. Gluck and Glenn Gunzelmann
29. Modeling and Formal Analysis of Human-Machine Interaction 433
Asaf Degani, Michael Heymann, and Michael Shafto
30. Queuing and Network Models 449
Yili Liu
31. Bayesian and Signal Detection Models 465
Jason S. McCarley and Aaron S. Benjamin
32. Judgment Analysis 476
Alex Kirlik
33. Modeling Decision Heuristics 490
Konstantinos V. Katsikopoulos and Gerd Gigerenzer
34. Establishing the Micro-to-Macro Link in Cognitive Engineering: Multilevel Models of Socio-Computer Interaction 501
Wai-Tat Fu and Peter Pirolli

Part Five • Cognitive Technologies in Engineered Systems

35. Configural and Pictorial Displays 517
Kevin B. Bennett and John M. Flach
36. Emergence in Organizations and Human Collective Intelligence 534
Stephen J. Guastello
37. Multimodal Displays: Conceptual Basis, Design Guidance, and Research Needs 556
Nadine B. Sarter
38. Ecological Interfaces 566
Catherine M. Burns
39. Uncertainty Visualization and Related Techniques 579
Ann M. Bisantz
40. Adaptive Automation 594
David B. Kaber
41. Distributed Communities of Practice 610
Anna T. Cianciolo and Karen M. Evans
- Index 623

This page intentionally left blank

PART

1

Cognitive Engineering:
History and
Foundations

This page intentionally left blank

Introduction to the Handbook

John D. Lee *and* Alex Kirlik

Abstract

Since its origins some 30 years ago as a subdiscipline of both human factors and cognitive science, cognitive engineering has grown into a diverse yet coherent body of research and practitioner activity focused on informing the human-centered design of engineered systems and workplaces. To introduce this handbook, we first provide a brief statement of the perspectives that gave rise to the selection and organization of the research it presents. We then situate cognitive engineering historically, both as a maturation of cognitive science to embrace applications and as an outgrowth and extension of human factors enabled and required by developments in information technology and automation. Finally, in the concluding sections of the chapter, we apply a cognitive engineering approach to the book itself, using large-scale data collection and analysis, statistical modeling, and pictorial visualization to provide the reader with a set of windows into the contents of this handbook. We hope that these windows function as effective user-centered aids to facilitate both efficient use and to communicate the research themes comprising contemporary cognitive engineering to a broad audience of students, researchers, and practitioners.

Key words: historical foundations, themes, text analysis, visualization

Handbook Contents and Organization

Cognitive engineering is an interdisciplinary approach to the analysis, modeling, and design of engineered systems or workplaces, especially those in which humans and automation jointly operate to achieve system goals. The field emerged in the 1980s in response to the increased complexity of the challenges faced by system designers and by the enhanced array of opportunities afforded both by these technologies and by the maturation of cognitive science, better enabling it to inform design.

A brief inspection of the table of contents shows that this handbook is organized into four major sections:

- Cognition in Engineered Systems
- Cognitive Engineering Methods
- Cognitive Engineering Models
- Cognitive Technologies in Engineered Systems

For the reader with at least some acquaintance with cognitive engineering research, the most noteworthy aspect of this organization is most apparent when contrasted with an alternative framework organized instead around application domains:

- Cognitive Engineering in Health Care
- Cognitive Engineering in Aviation
- Cognitive Engineering in Highway Transportation
- And so forth

There are at least two reasons for being explicit in bringing the domain-general, as opposed to domain-specific, organization chosen for this handbook to the readers' attention.

First, while it is indeed true that most cognitive engineering projects require the input of substantial domain knowledge or expertise, domain specificity (of theories, models, best practices, design

approaches, etc.) has the potential to stand in the way of informative and efficient cross-domain generalization and the development of cognitive engineering into a mature engineering discipline in its own right. We were well aware of the extent to which some of this handbook's contributors would be content or pleased to identify their research with one application domain or another, as they may have found it professionally reinforcing to identify with an application domain in these times of intensive academic specialization. We were nevertheless delighted to find such a substantial base of cognitive engineering researchers willing to take on the challenge we posed to write their chapters in domain-general terms (that is, every author represented in this handbook).

Second, prior to our many dialogues with contributors to help identify suitable cross-domain topics for their contributions, and seeing the resulting manuscripts, we ourselves were unsure what themes and topic areas would emerge and sustain as viable, domain-general themes around which this handbook could be organized. Unsure if our requests would yield productive results, we nevertheless put the challenge to the cognitive engineering research community to communicate useful information to readers in a way that did not presume prior knowledge of the reader's domain of interest. We found that the community, in our judgment, succeeded admirably. On the basis of these results, we feel confident that this handbook will be of interest and use to a broad audience of practitioners, many of them engineers and computer scientists, involved with designing human-technology systems for a broad array of application domains. We also hope that the benefits to the student and academic communities will be similarly direct. This handbook may aid in selecting research problems that already show strong promise of domain generality, and thus broad relevance and impact, because current research has yet to explore and probe the full range and depth of issues involved.

Historical Foundations

We have already mentioned that one of the historical developments that led to the emergence of cognitive engineering is the maturation of cognitive science into a discipline whose theories, models, and methods are capable of guiding application. These developments have come mainly from two directions. The first involved the groundbreaking work toward creating computational models of cognition, initially, in the domain of human-computer

interaction (Card, Moran, & Newell, 1983). This research has proven seminal in prompting numerous extensions resulting in various “cognitive architectures” and related approaches to modeling cognitive performance in technological interaction. A variety of these approaches are described in the Cognitive Engineering Models section of this handbook, as well as in research volumes devoted to this approach (e.g., Gluck & Pew, 2005). Research such as this continues a longstanding appreciation for the fundamental role played by modeling in the analysis and design of human-technology systems (Elkind, Card Hochberg, & Huey, 1990; Rouse, 1980; Sheridan & Ferrell, 1974). Research methods grounded in modeling, whether quantitative, computational, or otherwise, are a hallmark of both professional and research activity in engineering. Cognitive engineering is not likely to be different.

A second route by which cognitive science matured into application in a manner that helped spawn the field of cognitive engineering is through the research of Donald A. Norman, first outlined in his chapter titled “Cognitive Engineering” in the 1983 volume *User Centered Systems Design*. Here, Norman laid out his influential and intuitive conception of the barriers to good design lying in the “gulf of execution” (how do I get it to work?) and “gulf of evaluation” (is it working as I intended?). It is useful to consider how Norman himself understood the nature of the discipline he was putting forward at the time:

Cognitive Engineering, a term invented to reflect the enterprise I find myself engaged in: neither Cognitive Psychology, nor Cognitive Science, nor Human Factors. It is a type of applied Cognitive Science, trying to apply what is known from science to the design and construction of machines. It is a surprising business. On the one hand, there is quite a lot known in Cognitive Science that can be applied. On the other hand, our lack of knowledge is appalling. (Norman, 1983, p. 31)

Norman's comments prompt one to consider why he made a contrast between what he was promoting and the much older discipline of human factors, which by almost any definition involves applying “what is known from science to the design and construction of machines” (e.g., Wickens, Lee, Liu, & Gordon-Becker, 2004). Norman believed there was a gap between what the discipline of human factors was offering at the time and what was needed to provide sufficient guidance for the

design of interactive technologies. Many of these gaps have developed as important research themes in this handbook and the field of cognitive engineering more generally.

Why has cognitive engineering emerged as either a separate or subdiscipline of human factors? Many perspectives on this question exist. Google Scholar provides one useful path forward. As of the time of this writing, the only publication with more Google Scholar hits using the search term “cognitive engineering” than Norman’s previously cited (1983) chapter (over 1600 citations) is Jens Rasmussen’s book *Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering*, with over 2300 citations. This monograph and Rasmussen’s classic (1983) article “Skills, Rules, and Knowledge; Signals, Signs, and Symbols, and Other Distinctions in Human Performance Models” are taken by many in the cognitive engineering community to be seminal publications and landmarks at the origin of the field.

In both his 1983 article and 1986 monograph, Rasmussen, a control engineer working to ensure the safety of nuclear power plants and operations, observed that, in the crucially important area of interface design, semantics had overtaken syntax as the chief barrier to effective system control and problem diagnosis. It was not that system operators had great difficulty perceiving or attending to proximal information displays, but rather in understanding what they meant. This observation implied that the lion’s share of human factors knowledge on how to present information at an interface to best support perception or attention was, while necessary, far from sufficient to ensure effective human-machine interaction mediated by interface displays. Displaying information to which an operator can attend and perceive was important, but to Rasmussen, design guidance of this kind was insufficient if it did not also foster operator comprehension or understanding.

Rasmussen understood meaning in terms of external reference. The operator’s ultimate task, Rasmussen noted, is to monitor, control, and diagnose (and so forth) a plant or technology “behind” the interface, so to speak. The operator’s task is not merely to attend, perceive, and manipulate the proximal interface itself, although interface manipulation skills have historically served as the object of study for the lion’s share of traditional human factors research. Instead, for Rasmussen, the proximal interface must be considered functionally, not as the ultimate or end target of human interaction,

but instead as a window to a distal plant or remote environment comprising the true target of work. Just as Bruner (1973) had characterized cognition as “going beyond the information given,” Rasmussen (1983) described an operator’s cognitive task in terms of exactly the same sort of going beyond, but in this case, going beyond the given interface.

This characterization applies not solely to process control, but equally to modern “knowledge workers” (Zuboff, 1984) more generally, whose windows to the world of work increasingly consist of computer interfaces of one sort or another and, as such, who are rarely able to perceive and manipulate the objects of their work in a direct fashion. Additionally, the heightened emphasis given to knowledge-based behavior in Rasmussen’s research and in an array of related cognitive science research on expertise (e.g., de Groot, 1978; Simon & Chase, 1973) served as one impetus to a line of research focused on the nature of expert decision making in cognitive engineering contexts (Klein, 1989) and the mechanisms associated with human error by otherwise well-trained, well-motivated human operators (Reason, 1990; Senders & Moray, 1991). Cognitive engineering’s primary focus on expert or otherwise knowledgeable humans is evident throughout this handbook. This focus is yet another factor that marks the discipline off from much, but not all, traditional engineering psychology and human factors research, which has historically focused on the behavior of humans with perhaps hours, but rarely months or years, of training and experience.

Rasmussen’s (1983) paper presented a conceptual framework that both acknowledged the importance of the large body of research that had grown up around relatively simple technological contexts in which the primary goal was safe and efficient interaction with a proximal interface or workplace, yet nevertheless indicated a need for a novel theory and method for better understanding the cognitive activities of knowledge workers. Rasmussen’s observations have proven prescient: The research problems that occupy the lion’s share of the attention of today’s cognitive engineers are those in which technology is not viewed as the end target of human interaction, but rather as an intermediary through which humans interact with the actual objects of work.

It is also worthwhile to consider how Rasmussen (1983) laid out what he believed to be necessary for cognitive engineering to meet its goals:

In our work, concern is with the timely development of models of human performance which can be useful for the design and evaluation of new interface systems. For this purpose, we do not need a single integrated quantitative model of human performance but rather an overall qualitative model which allows us to match categories of performance to types of situations. In addition, we need a number of more detailed and preferably quantitative models which represent selected human functions and limiting properties within the categories. The role of the qualitative model will generally be to guide overall design of the structure of the system including, for example, a set of display formats, while selective, quantitative models can be used to optimize the detailed designs. (p. 264)

Some 30 years after Rasmussen stated these objectives for future research, we expect that the reader will see, as illustrated by this handbook, that the array of contemporary cognitive engineering products consists largely of a toolbox of conceptual or qualitative frameworks together with a set of more formal techniques and quantitative models for detailed performance prediction.

Rasmussen's research was also influential in providing cognitive engineering's orientation to a unit of analysis consisting of a human-technology system, or perhaps even a human-technology-environment system, rather than the human in isolation. This was not a new idea within the engineering-oriented, human performance modeling tradition (see Pew, 2008, for a historical overview), yet Rasmussen's observations on the fundamental *ecological* nature of cognitive engineering resonated with researchers interested in grounding the psychology of cognitive engineering in a scientific footing other than solely information processing theory. Cognitive engineering researchers such as Vicente (Vicente & Rasmussen, 1990; Vicente, 1999); Woods and Hollnagel (Hollnagel, Mancini, & Woods, 1988; Woods & Hollnagel, 2006); and Flach (1990) have each pursued cognitive engineering approaches influenced by the ecological theory of perceptual psychologist James J. Gibson, or, more importantly, on a unit of analysis spanning the human, cognitive tools, and the work environment. Along similar lines, and though grounded in a computational rather than an ecological framework, the pioneering research of Hutchins (1995) and his colleagues (Hollan, Hutchins, & Kirsh, 2000) on distributed cognition also brings to cognitive engineering an approach that seeks to account for how cognitive

resources both internal and external to the human might combine to enable the types of performance observed in technological systems. At a general level, the guiding theoretical orientation behind all these approaches is that cognitive engineering concerns the analysis and design of integrated, human-technology systems. This general orientation is evident throughout this handbook.

Another research theme central to this handbook concerns the challenge of achieving a safe and productive coupling of humans and automation. To the extent that the discussion above has been useful in communicating the pioneering influence of Jens Rasmussen's research on cognitive engineering, the research of Thomas B. Sheridan has played a similarly pioneering role in bringing the issues involved with human-automation interaction to the forefront of cognitive engineering research (Sheridan & Johannsen, 1976; Sheridan, 1992; Sheridan, 2002). Although much of the impetus for Rasmussen's research came from his observations of power plant technicians engaged in troubleshooting tasks (see Vicente, 2001, for a detailed history and overview), Rasmussen was also strongly influenced by the seminal research of Thomas Sheridan, who was actively engaged in the problems of remote control and monitoring of distant vehicles in contexts such as space and undersea exploration. Sheridan coined the term "supervisory control" to describe the situation in which the human is not in direct, manual control of a system or process, but instead inputs commands to automated systems that themselves act directly on the distal system, process, or vehicle.

This seminal research by Sheridan has spawned dozens of studies over the past decades trying to characterize human-automation interaction with models or taxonomies, to understand the consequences of introducing automation into systems or workplaces, to identify and describe human tendencies in dealing with those consequences, and to identify design principles, frameworks, and techniques to support human operators or workers in doing so. As will be seen in the following sections of this chapter, if one had to name a single key topic central to this handbook, the impact of automation and information technology on the human's role in engineered systems would be that key topic.

Finally, to both close this section on historical foundations and to set the stage for a discussion of the handbook's contents in detail, it should be mentioned that cognitive engineering has broadened its focus even further in recent years to include a consideration of how teams and organizations

communicate and collaborate in the performance of cognitive tasks (Salas & Fiore, 2004). While hardly a new idea, the proliferation of information and communication technologies that increasingly mediate what was once direct human interaction has highlighted the importance of these social factors. A variety of chapters in this handbook—including those on communication, teamwork, conducting experiments with teams, and the design of organizations and communities of practice—illustrate this rapidly growing dimension of cognitive engineering research.

Cognitive Engineering Themes and the Handbook Contents

These historical themes are reflected in the chapter structure and in the associated content. Figure A.1 shows a word cloud based on the contents of the 41 chapters, which provides a simple visualization of the handbook contents. The size of each word is proportional to its frequency of occurrence in the book—the large words occur often. This representation suggests the broad scope of cognitive engineering, spanning the individual operator to teams and organizations, with a focus on how systems of people and technology, often in the form of automation, influence performance. This word cloud provides a holistic view of the handbook contents that can be challenging to extract from reading the individual chapters. While useful, the word

cloud is limited because it provides no link to back to the chapters. A reader seeing the importance of models and systems from the word cloud would not know what chapters to read to learn more about these topics. Formal analysis of the text represented in the word cloud can help readers to navigate the complex field of cognitive engineering.

The 41 chapters contained in this handbook demonstrate the diversity of perspectives that define the field of cognitive engineering. The organization of these chapters in the handbook represents one way of compiling this content, and we have made a concerted effort to structure these chapters in a logical fashion by placing related material together. At the same time, reading through the chapters or scanning the table of contents might not convey an adequate understanding of the field and might not lead a reader to a set of chapters of particular interest. Not all readers will want to read all 41 chapters, and readers will approach the handbook with diverse backgrounds and objectives. Different readers will need different tables of contents to satisfy their needs. We apply text analysis to the chapters to identify common themes and connections between the chapters that a single table of contents cannot provide.

We hope these themes and connections will support a focused and individualized reading of the handbook that meets the particular needs of each reader. As an example, a designer interested in situation awareness could start by reading chapters that contain “situation awareness” in the title, but the handbook is not organized to identify related chapters that might also be of interest. There is no “situation awareness” section in the handbook. The titles of other chapters might not reveal their relevance, and chapters located before or after the chapters with “situation awareness” in their title might not be particularly closely related to situation awareness. A reader interested in situation awareness might then be left to search the index. Many readers will approach the handbook with similarly individual perspectives. To support readers who desire for a focused reading of the handbook, this chapter provides a systematic analysis and representation of the handbook contents.

Three analyses support a more focused and individualized reading of the handbook. First, we identify groups of similar chapters based on the relative frequency of words occurring in each chapter. One might think of this as identifying clusters of chapters that have similar word clouds. Second, we describe topics contained in these chapters. Even chapters




Figure A.1 A word cloud of the handbook content, highlighting the modeling and systems perspectives in understanding the influence of information technology and automation.

that fall into the same cluster might address different topics, and so the topics contained in each chapter indicate why a chapter belongs to a particular cluster and also indicate chapters that share the same topic even if it belongs to a different cluster. Third, we describe how shared topics connect chapters into a network of chapters. This network highlights chapters particularly central to the field as those chapters that contain themes that are shared by many other chapters. In combination, these analyses provide an alternate table of contents to the handbook that we hope will help readers navigate the field of cognitive engineering.

The text analysis techniques applied to the chapters in this book are based on the term frequency data represented in Figure A.1. Each chapter is reduced to a vector that tabulates the frequency of each word used in the chapter. The handbook can then be represented as a matrix, with each chapter as a row and each column representing the frequency of occurrence of words, such as those shown in Figure A.1. The relative frequency of occurrence of words across chapters can be analyzed as numeric data using techniques such as cluster analysis. This “bag of words” approach to text analysis does not include any information regarding the meaning of particular words or their relationship to each other within sentences. Even so, analysis of such term-frequency data often provides a surprisingly insightful view into the concepts contained in a set of documents (Landauer & Dumais, 1997; Deerwester, Dumais, Furnas, Landauer, & Harshman, 1990).

The statistical package R 2.14.1 (R Development Core Team, 2011) supported the text analysis of the handbook chapters, with text mining packages *tm* (Feinerer, Hornik, & Meyer, 2008) and *topicmodels* Grün & Hornik (2011). The graphics packages *ggplot2* (Wickham, 2010) and *igraph* (Csárdi & Nepusz, 2006) were used to visualize the results.

Applying Ward’s method of hierarchical clustering to the data from the term-document matrix identifies similar chapters. Here, chapter similarity is based on the Euclidian distance between chapters defined by the relative frequency of each term contained in each chapter. Documents that use the same terms with the same relative frequency will be close to each other and so will fall into the same cluster. Figure A.2 shows the hierarchical cluster analysis, with the top of the hierarchy showing two sets of chapters and the bottom of the hierarchy showing individual chapters. A cut point midway in the hierarchy produces 13 clusters of chapters. The chapters generally cluster according to the table of contents

of the book. Many clusters include chapters from a similar section of the book, such as chapters III.1, III.2, and III.3, 14, 15 on task analysis, work analysis, and decision-centered design. Others that are not co-located in the book are strongly related, such as chapters II.5 and III.4 on situation awareness and chapters II.4 and IV.6 on judgment. Considering how these clusters combine in the hierarchy shows that the clusters of clusters also correspond to the grouping in the book. The cluster that combines the three clusters beginning with the third cluster from the left is almost exclusively composed of chapters from section III. Importantly, this analysis uses only word frequency, with no reference to the chapter structure, and yet it the clusters reflect the structure of the book surprisingly closely.

Topic analysis reveals the themes contained in the chapters that influence cluster membership. Based on a latent Dirichlet allocation approach (Grün & Hornik, 2011), the text of the 41 handbook chapters reveals 22 distinct topics. These topics can be represented by word clouds of the terms most important in defining each topic. Figure A.3 shows the word clouds of the 22 topics, and from these word clouds of the topics names were defined. These topics and their associated word clouds provide a much richer description of the handbook content than a single word cloud in Figure A.1. These word clouds describe some of the common themes of the field of cognitive engineering, but like the overall word cloud in Figure A.1, by themselves they do not direct readers to particular chapters.

Figure A.4 shows that each topic occurred in at least three chapters and that a few topics occur in almost one quarter of the chapters. The topics of “decision heuristics,” “simulation for safety,” and “communication” each occur in four or fewer chapters, whereas the topics of “team coordination” and “practice and learning” occur in at least nine chapters. Figure A.5 shows the distribution of these topics across the chapters. Chapters are listed vertically, grouped according to the cluster analysis. The dark and light grays differentiate neighboring clusters. The horizontal axis indicates the topics of the handbook. Three topics describe each chapter, and the size of the circle represents which topic is the primary, secondary, or tertiary topic of each chapter—the large circle represents the primary topic. Generally, chapters in the same cluster share topics—the chapters on judgment share the primary topic of judgment. Chapters in some clusters do not share topics as uniformly, as in the case of the




Figure A.2 Ward’s method for hierarchical cluster analysis identifies 13 clusters of chapters using Euclidian distance of each chapter.

chapter on trust and the chapter on communities of practice. These chapters belong to the same cluster but do not share any topics. The vertical lines highlight the topics of “team coordination” and “practice and learning,” which occur in the greatest number of chapters.

Figure A.5 provides a valuable tool for understanding the content of the handbook by highlighting clusters of chapters and also by highlighting particular topics covered by these chapters. The combination of chapter clusters and the topics can help readers identify particular sets of chapters that

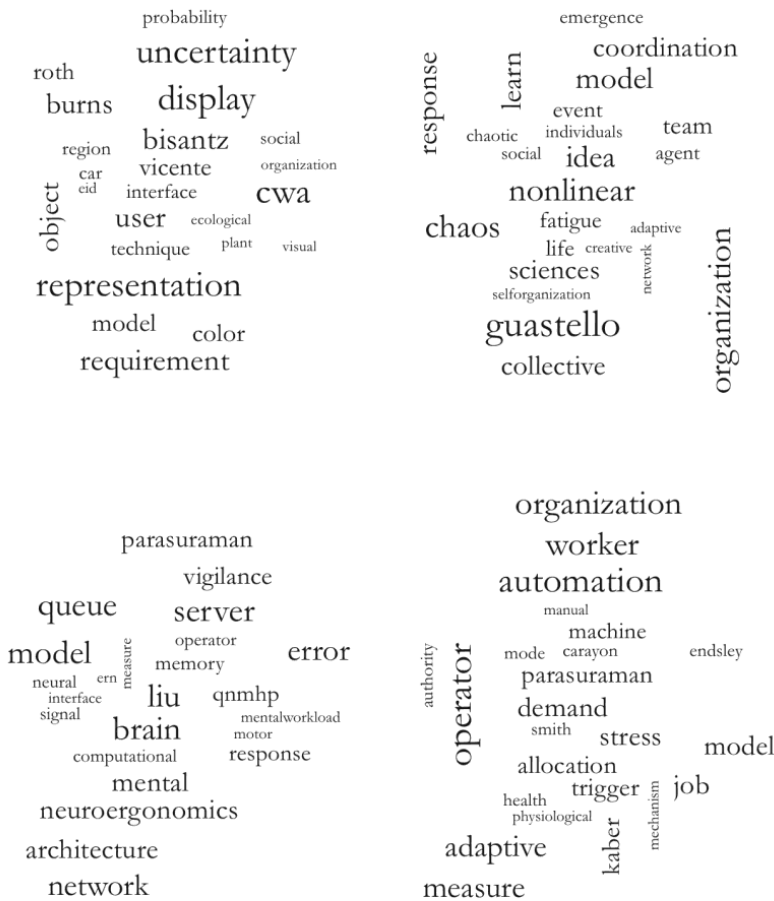
address topics of interest that might not be obvious from the chapter titles or by the structure of the table of contents. Locating the topic on the horizontal axis and then tracing that column upward to find the circles that indicate chapters that include a particular topic can identify a set of chapters addressing an issue that might not be clear from the chapter's title or section of the table of contents. Figure A.5 can act as an alternate table of contents that makes it possible for readers to quickly identify chapters that are most likely to meet their needs.

Figure A.5 shows that many chapters share topics with other chapters. These shared topics can be considered as links between chapters that form a network of chapters. In this network, a chapter might be connected to one or two other chapters or to many chapters. The structure of this network based on shared topics places each chapter into a rich context of connections with other chapters.

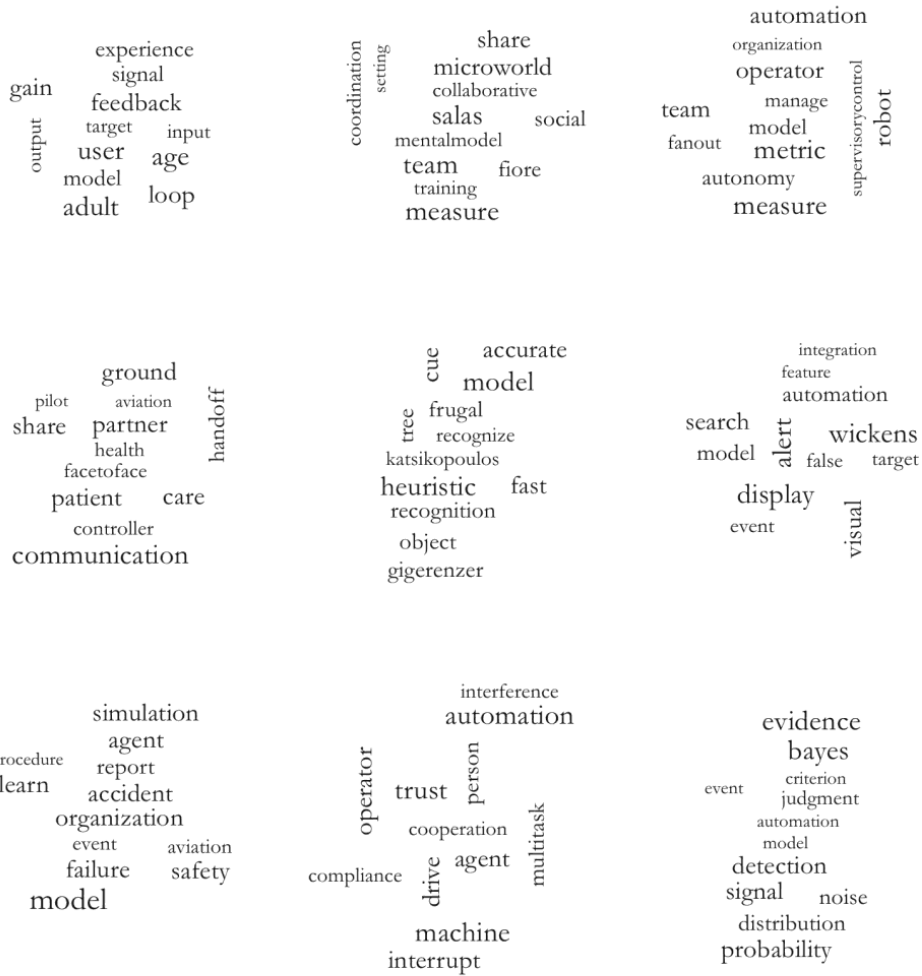
Network analysis measures provide a way to quantify features of the chapter network defined

by shared topics. One network analysis measure considers the frequency of links between nodes in the network—degree. Nodes with many connections to other nodes have high degree or centrality. In the network of chapters, highly central nodes are those that share topics with many other chapters. Another network analysis measure—betweenness—considers the number of links to nodes that are linked to many others. Such nodes have a short path to other nodes in the network. In the network of chapters, chapters with high betweenness are linked to chapters that touch on many of the topics that define cognitive engineering. Multidimensional scaling uses the links between nodes to place nodes that share similar patterns of connections near each other in a two-dimensional space.

Figure A.6 shows the network structure of the handbook defined by shared topics. Highly central chapters are indicated by large nodes, and chapters with high betweenness are indicated by



Topics 1–4: Uncertainty representation, collective coordination, neural architectures, automation and organizations
Figure A.3 Word clouds of each of the 22 topics



Topics 5–13: Feedback and control, team coordination, supervisory control, communication, decision heuristics, directed attention, modeling safety, trust in automation, probabilistic models

Figure A.3 (continued)

nodes that are dark red. In this network, chapters that are highly central also tend to have high betweenness.

Readers can use space represented by Figure A.6 to find chapters in the same “neighborhood” to identify a set of chapters of interest. For example, the chapter V.3 on ecological interface design, in the lower right of the network, is in the neighborhood of chapter III.2, which addresses a highly related topic of cognitive work analysis, and relatively close to chapters on uncertainty visualization and configurational displays. Similarly, chapters on queuing models of cognition, adaptive automation, formal models of automation, attention, and neuroergonomics are all closely clustered in the middle left of the figure. Finding a chapter of interest in this space and then surveying its neighbors can guide a focused exploration of the handbook.

More generally, the space defined by the network in Figure A.6 reflects the broad themes of the historical development of cognitive engineering. The centrality and prominence of team cognition is perhaps the most notable feature of this space. Slightly above these team chapters is a series addressing supervisory control and human-automation interaction. This position suggests an important trend of technology to share many of the same features of a human team member, blurring the distinctions of supervisory control of humans versus that of automation. From this center, to the right, many chapters address cognitive engineering in broader organizations and communities, pointing to the need to consider the engineering of organizations and social networks. To the left of the teamwork chapters, many chapters focus on individual cognition, addressing topics of attention, decision making, and multitasking. The horizontal



Topics 14–22: Expertise, social models, simulation for safety, judgment, supporting decisions, machine models, representations and artifacts, practice and learning, computational models

Figure A.3 (continued)

dimension broadly moves from individual cognition to the cognition of teams, and then networks.

In contrast to the individual-to-network span of the horizontal axis, the vertical axis broadly contrasts between cognitive task analysis at the top of the figure and cognitive work analysis at the bottom of the figure. A focus on the ecology and an analysis of the constraints and dynamics of the engineered system anchor the vertical dimension at the bottom of Figure A.6—cognitive work analysis. The top of the figure focuses on the cognitive processes and information needed to support decisions and situation awareness—cognitive task analysis. In some sense, the top of the figure represents the *cognitive* in cognitive engineering, and the bottom of the figure represents the *engineering*. The top of the figure represents the application of principles of cognitive psychology, and

the bottom represents the application of an engineering perspective to characterize the plant or technology “behind” the interface. Such a simple caricature of the chapters and cognitive engineering as a whole certainly glosses over many important distinctions, but does highlight important themes that have guided the profession since its inception.

Figure A.6 collapses a complex, multidimensional space to two dimensions. Likewise, the topic analysis collapses many subtle points to 22 topics. Like any model, such simplifications distort and fail to completely capture the phenomena of interest. Careful inspection of Figures A.5 and A.6 reveals instances where the model might fail to represent reality well. For example, given the broad dimensions of Figure A.6—with analysis of individuals on the left, teams in the center, and networks on the




Figure A.4 Frequency of occurrence of topics across chapters

right—one might expect the chapter on organization design to anchor the horizontal dimension on the right. One explanation for its location might stem from how terms are treated in the analysis, particularly the term “operator,” which is used to refer to a person and as a cognitive process in the chapter on queuing models. The text analysis techniques we used are blind to this distinction, highlighting the limits of any simple model of a complex field.

Nevertheless, we hope the necessarily imperfect representations in Figures A.5 and A.6 are useful tools for exploring this handbook and the field of cognitive engineering.

Conclusions

The field of cognitive engineering and this handbook are large and complicated. Consequently, no single chapter structure will serve all readers’ needs.




Figure A.5 Chapters, chapter clusters, and the associated topics. The vertical lines highlight the two most frequently occurring topics.

Text analysis reveals clusters of chapters, topics, and connections between chapters based on these topics that provide alternate ways of exploring the material in the handbook and understanding the field. The forces

that initiated the emergence of cognitive engineering nearly 30 years ago have not abated, but instead have intensified. Designers face increasingly complex technologies that, properly engineered, enable people to



Figure A.6 A network representation of handbook chapters with links defined by shared topics arrayed according to multidimensional scaling.

work with and through such technology in new and increasingly productive ways. The chapters in this handbook profile advances and remaining challenges in designing for cognitive work.

References

- Bruner, J. (1973). *Going beyond the information given*. New York, NY: Norton.
- Card, S. K., Moran, T. P., & Newell, A. (1983). *The psychology of human-computer interaction*. Hillsdale, NJ: Erlbaum.
- Csárdi, G., & Nepusz, T. (2006). The igraph software package for complex network research. *InterJournal Complex Systems, CX.18*, 1695.
- Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., & Harshman, R. (1990). Indexing by Latent Semantic Analysis. *Journal of the American Society for Information Science and Technology, 41*(6), 391–407.
- de Groot, A. D. (1978). *Thought and choice in chess* (2nd ed.). Paris, France: Mouton De Gruyter.
- Elkind, J. I., Card, S. K., Hochberg, J., & Huey, B. M. (1990). *Human performance models for computer-aided engineering*. New York, NY: Academic Press.
- Flach, J. M. (1990). The ecology of human-machine systems I: Introduction. *Ecological Psychology, 2*(3), 191–205.
- Gluck, K. A., & Pew, R. W. (2005). *Modeling human behavior with integrated cognitive architectures: Comparison, evaluation, and validation*. Mahwah, NJ: Erlbaum.
- Grün, B., & Hornik, K. (2011). topicmodels: An R package for fitting topic models. *Journal of Statistical Software, 40*(13), 1–30.
- Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction, 7*(2), 174–196.
- Hollnagel, E., Mancini, G., & Woods, D. D. (1988). *Cognitive engineering in complex, dynamic worlds*. New York, NY: Academic Press.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.

- Klein, G. (1989). Recognition-primed decisions. In W. B. Rouse (Ed.), *Advances in man-machine systems research* (pp. 47–92). Greenwich, CT: JAI Press.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*(2), 211–240.
- Norman, D. A. (1983). Cognitive engineering. In D.A. Norman & S.W. Draper (Eds.), *User centered systems design* (pp. 31–62). Hillsdale, NJ: Erlbaum.
- Pew, R. W. (2007). Some history of human performance modeling. In W. Gray (Ed.), *Integrated models of cognitive systems* (pp. 29–44). New York, NY: Oxford University Press.
- Pew, R. W. (2008). More than 50 years of history and accomplishments in human performance model development. *Human Factors*, *50*(3), 489–496.
- R Core Team, R. C. (2012). R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing.
- Rasmussen, J. (1983). Skills, rules, and knowledge; Signals, signs, and symbols, and other distinctions in human performance models. *IEEE Transactions on Systems, Man, and Cybernetics*, *SMC-13*(3), 257–266.
- Rasmussen, J. (1986). *Information processing and human-machine interaction: An approach to cognitive engineering* (North-Holland Series in Systems Science and Engineering, No. 12). Elsevier.
- Reason, J. (1990). *Human error*. Cambridge, England: Cambridge University Press.
- Rouse, W. B. (1980). *Systems engineering models of human-machine interaction*. Amsterdam, The Netherlands: North-Holland.
- Salas, E., & Fiore, S. (2004). *Team cognition*. Washington, DC: APA Press.
- Senders, J., Moray, N. P. (1991). *Human error: Cause, prediction, and reduction*. Series in applied psychology. Hillsdale, NJ: Erlbaum.
- Sheridan, T. B. (1992). *Telerobotics, automation, and supervisory control*. Cambridge, MA: MIT Press.
- Sheridan, T. B. (2002). *Humans and automation: System design and research issues*. New York, NY: John Wiley & Sons.
- Sheridan, T. B., & Ferrell, W. R. (1974). *Man-machine systems*. Cambridge, MA: MIT Press.
- Sheridan, T. B., & Johannsen, J. (1976). *Monitoring behavior and supervisory control* (NATO special program panel on human factors). New York, NY: Plenum Press.
- Simon, H. A., & Chase, W. G. (1973). Skill in chess. *American Scientist*, *61*, 394–403.
- Vicente, K. J. (1999). *Cognitive work analysis: Toward safe, productive and healthy computer-based work*. Boca Raton, FL: CRC Press.
- Vicente, K. J. (2001). Cognitive engineering research at Riso from 1962–1979. In E. Salas (Ed.), *Advances in human performance and cognitive engineering research* (Vol. 1, pp. 1–58). Oxford, England: Elsevier.
- Vicente, K. J., & Rasmussen, J. (1990). The ecology of human-machine systems II: Mediating “direct perception” in complex work domains. *Ecological Psychology*, *2*(3), 207–249.
- Wickens, C. D., Lee, J. D., Liu, Y., & Gordon-Becker, S. (2004). *An introduction to human factors engineering* (2nd ed.). Upper Saddle River, NJ: Prentice-Hall.
- Wickham, H. (2010). A layered grammar of graphics. *Journal of Computational and Graphical Statistics*, *19*(1), 3–28.
- Woods, D. D., & Hollnagel, E. (2006). *Joint cognitive systems: Patterns in cognitive systems engineering*. Boca Raton, FL: CRC Press.
- Zuboff, S. (1984). *In the age of the smart machine*. New York, NY: Basic Books.

PART 2

Cognition in
Engineered Systems

This page intentionally left blank

The Closed-Loop Dynamics of Cognitive Work

John M. Flach, Kevin B. Bennett, Richard J. Jagacinski, Max Mulder, and Rene van Paassen

Abstract

This chapter provides a tutorial introduction to the logic of closed-loop systems. A series of examples is presented to illustrate the dynamics of closed-loop systems and to contrast the behavior of these systems with expectations suggested by the logic of open-loop causal systems. In particular, the examples show that “feedback” per se is insufficient to guarantee convergence on a target. The chapter explores some of the issues that must be addressed in order to determine whether closed-loop systems will be stable. Parallels to the phenomena of cognition in the wild are highlighted, and the case is made that the logic of closed-loop systems is an essential foundation for understanding the behavior of cognitive systems.

Key Words: adaptive control, abduction, closed-loop control, observer, self-organization

Introduction

In 1942, at a conference sponsored by the Macey Foundation to promote interdisciplinary discussions about neuroscience, Arturo Rosenblueth introduced a radical challenge to classical ideas of cause and effect. He introduced the construct of “circular causality.” This new construct was motivated by collaborations with Norbert Wiener and Julian Bigelow to understand stability in feedback control systems (Rosenblueth, Wiener, & Bigelow, 1943; Wiener, 1948/1961). The neuro- and social scientists at the meeting (including Warren McCulloch, Gregory Bateson, and Margaret Meade) immediately identified this idea as relevant to their work—easily identifying concrete examples of circular coupling in the biological and social systems that they were studying (Conway & Siegelman, 2005).

There is no doubt that the “cybernetic hypothesis” that emerged from the work of Norbert Wiener and his colleagues had an enormous impact on the trajectory of cognitive science. For example, feedback was a central theme in Miller, Galanter, and

Pribram’s (1960) influential book *Plans and the Structure of Behavior*.

The general pattern of reflex action, therefore, is to test the input energies against some criteria established in the organism, to respond if the result of the test is to show an incongruity, and to continue to respond until the incongruity vanishes, at which time the reflex is terminated. Thus, there is “feedback” from the result of the action to the testing phase, and we are confronted by a **recursive loop**. (p. 26, emphasis added)

Miller et al. introduced the TOTE unit (Test-Operate-Test-Exit) as a simple process to illustrate the concept of feedback using the example of hammering a nail. They conclude:

It may seem slightly absurd to analyze the motions involved in hammering a nail in this explicit way, but it is better to amuse a reader than to confuse him. It is merely an illustration of how several simple TOTE units, each with its own test-operate-test **loop**, can be embedded in the operational phase of a larger unit

with its particular test-operate-test **loop**. Without such an explicit illustration it might not have been immediately obvious how these **circles within circles** could yield hierarchical trees. (p. 37, emphasis added)

Miller et al. used the “feedback” concept as a means to motivate their exploration of plan hierarchies, and in turn they helped to motivate a 50-year program of research to explore human information processing. However, with some notable exceptions, the cybernetic hypothesis has generally been trivialized, and the field of cognitive science as a whole has failed to appreciate the dynamics of circles and loops. Most in the field have persisted in viewing the phenomenon of human information processing through the lens of the logic of simple open-loop causal systems. Most have failed to understand the full implications of the concept of circular causality.

The classical image of an information processing system is an open-loop series of “stages.” These stages form a chain from the stimulus to the response, which is treated as a string of dominoes with the “response” from one stage being the “stimulus” for the next in a sequence of discrete acts. In this framework, there is an implicit arrow of time, such that some stages are seen as logically prior (as input or cause) to other stages. In fact, whole research programs are designed around single stages—as if each stage can be understood independently from the other stages.

However, if the dynamic of cognition is closed-loop, this image will not work. A better image is to think of the stages as being linked within a web—so that an impact anywhere within the web reverberates through the whole web. In a circular dynamic, each stage is simultaneously providing input to other stages and receiving input from those stages. There is no implicit arrow of time. Therefore, relative to the overall recursive, circular flow of influence, no stage can be simply categorized as either “cause” or “effect” relative to another stage. Also, note that the loop is closed through the ecology, so that the ecology is an intrinsic component of the web, simultaneously contributing as both cause and effect (consequence).

The goal of this chapter is to reintroduce the idea of circular causality—to provide a pedagogically sound but analytically simple introduction to the logic of feedback control. We hope to correct some common misconceptions about the nature of feedback and to help those interested in cognition in the wild to appreciate the dynamics of circles and

loops and the implications for studying situated cognition.

The Stability Problem

One of the most pervasive misconceptions about feedback is that the presence of feedback is sufficient to explain stable progression toward a goal. In fact, the central problem of control theory as a field of study is to identify the special conditions under which feedback will result in stable control. Stable progression to a goal is NOT guaranteed by the presence of feedback. In Miller et al.’s example of hitting the nail, there is an assumption that each swing is executed exactly right so that the only relevant feedback is whether the nail is flush or not. However, what if the nail position shifts from vertical or the nail bends? How will the swing need to be adjusted to achieve the goal of the carpenter? Or should the goal be shifted to removing the nail and starting over? Hammering a nail so that it satisfies the goals of the carpenter is not as trivial a problem as Miller et al.’s description suggests. They prematurely dismiss the “circles within circles” in order to get to hierarchies, leaving the dynamics of feedback systems unexplained; for the most part, cognitive science has persisted in ignoring the dynamic of circles.

To illustrate the stability problem, consider the simple controller illustrated in Figure 1.1. This system consists of three components: a time delay that reflects the time to process error feedback, a gain that determines the magnitude of correction as a result of the error at each moment, and an integrator that effectively sums or concatenates the corrections across moments to determine the next response. Let’s consider the time delay to be a fixed property of the system (e.g., representing the reaction time from seeing an error to initiating a corrective action). Thus, the lone free parameter is the gain, which can be treated as a simple multiplier of the error signal. As input to the system, we will use a step function. Tracking a step can be considered analogous to a target acquisition task. When the signal changes from one level (position) to another, this is analogous to the target appearing—requiring the person to move a cursor from one fixed (home) position to another fixed position (the target).

Can you guess what value of the gain will insure that the output of the system illustrated in Figure 1.1 will follow the target (i.e., so that the output will shift from the home position to reliably converge on the new target position commanded by the step input)? You might guess that a gain




Figure 1.1 A simple feedback system with a step input. How will the output response vary as a function of the gain parameter? Under what conditions will the output converge to match the input?

of 1 would be ideal, insuring a one-to-one match between the input and the output. This is not a bad guess, and depending on the magnitude of the time delay, this may well produce a satisfactory solution. However, it probably will not be optimal. What would happen if the gain were reduced to less than 1 or increased to more than 1? How would this affect the output? Would it cause an offset between the target and the output, such that the size of the output step would scale proportionally to the magnitude of the gain? This is what would happen if this were an open-loop system. However, in this closed-loop system the gain does not determine the size of the output signal, but rather the speed at which the output will converge toward the input. Very low gains will result in a “sluggish” response to the input. The output will eventually reach the target position, but it will take a long time to get there (Figure 1.2). As gain is increased (i.e., sensitivity to error is increased), the speed of correction will increase. However, at some point the higher gains will cause the system output to overshoot the target, before gradually converging back to the target. At even higher gains the output will oscillate around the target before settling down, and at still higher gains the oscillations will actually grow over time—so that the output diverges, never settling down on the target (Figure 1.2).

In general, the range of stable gains will depend on the size of the time delay. As time delays become larger, the range of gains that will produce stable control will get smaller. That is, with longer time delays the system will need to be less sensitive (i.e., lower gains—more conservative or more cautious) in responding to errors in order to avoid oscillatory or unstable responses (e.g., Jagacinski, 1977). In designing automatic control systems, much of the attention is given to eliminating unnecessary time delays and then determining an appropriate gain for the forward loop—in order to insure a fast but stable response to errors. In cases where long time delays are unavoidable (e.g., tele-operation over large distances), there may be no satisfactory gain for proportional control. In these cases, a discrete style of control (small adjustment, wait, small adjustment, wait, ...) will generally prove to be a more satisfactory means of control.

Figure 1.2 illustrates the range of output behaviors as a function of the gain parameter. It should be clear from this illustration that feedback alone does not guarantee that the output will converge to the target. This is one of the fundamental issues of control theory—to identify those special conditions that lead to stable control, or more generally to study those factors that determine the boundaries of stable control. In fact, it was the parallels between the behavior




Figure 1.2 This diagram illustrates a range of outputs for the control system illustrated in Figure 1.1 as a function of the gain parameter. For very low gains, the response will “sluggishly” converge to the target without overshoot. For very high gains, the response will become unstable with oscillations that diverge from the target. At intermediate gains, the response will range from relatively rapid convergence to the target without overshoot to converging oscillations as gain is increased within this range.

of unstable engineered control systems and behavior associated with some motor ataxias (purpose tremor) that stimulated the interdisciplinary collaboration between Wiener and Rosenbluth. These parallels suggested that the logic of circles was not simply an engineering problem, but a general problem of any system that adjusts its behavior based on feedback.

In natural closed-loop systems (with inevitable time delays) there will always be a speed-accuracy trade-off. That is, there will always be a limit to how fast a goal can be reliably approached (i.e., a limit to the gain). In human-machine systems, the gain is typically a joint function of the plant (e.g., the gain setting for the mouse, the gain of the control stick or steering wheel, and/or the vehicle dynamics) and the human (e.g., the scaling of the manual response to observed error). You may notice this when you use a computer system that has the mouse gain set differently than on your own computer. If the gain is higher, you may find yourself initially overcorrecting and oscillating around the target before finally capturing it. If the gain is lower, then the new system will feel “sluggish.” However, it will typically not take you long to recalibrate to the system gain and to make adjustments so that the combined gain (human + plant) results in reliable target acquisition.

Controls on high-performance aircraft typically have very high gains, so that pilots must compensate by making only small adjustments (lower gain) in response to errors so that the net gain (human + plant) will meet the constraints for stable control. It is not unusual to see novices lose control (i.e., generate pilot-induced oscillations) when they first try to fly a simulation of a high-performance aircraft. That is, their initial gain is too high! Research on human tracking shows that humans can learn to adjust their gain so that the human-vehicle system behavior is stable. Effectively, McRuer’s classic crossover model predicts that the effective control dynamics of the human-machine system will (with practice) behave much like the controller illustrated in Figure 1.1 (McRuer & Jex, 1967). That is, the combined system dynamics will approximate the behavior of a system with a time delay, gain, and integrator in the forward loop. This will be true for a fairly broad range of different plant (or vehicle) dynamics. In fact, the optimal control model predicts that for a range of vehicle dynamics humans can learn to conform fairly closely to the ideals prescribed by normative linear models for optimal tracking of randomly appearing input signals (Kleinman, Baron, & Levison, 1971). That is, the human will choose gains that minimize

a weighted combination of mean squared control speed (effort) and mean squared error (accuracy). Similar adaptations regarding plant dynamics, speed, and accuracy can be seen in the Fitts’ Law paradigm (e.g., Jagacinski & Flach, 2003).

There are several key points that we hope to have made in this section:

1. The intuitions derived from simple causal assumptions (based on open-loop dynamics) can be misleading when applied to closed-loop systems. That is, increasing gain determines the speed of convergence to the command input in a closed-loop system, NOT the relative magnitude or scale of the response as would be true in the case of an open-loop system.

2. All natural closed-loop systems will exhibit a speed-accuracy trade-off. That is, there will generally be a limited range of gains that result in satisfactory stable performance. When gain is too low, the response will be sluggish (too slow). When gain is too high, the response will become oscillatory, and with very high gain the oscillations can become unstable (increasingly diverging from the target).

3. The main point of this section is that *the presence of feedback alone does not insure stable convergence on a target or goal*. Thus, the presence of feedback is not the answer that explains behaviors guided by a goal. Rather, the logic of circles provides the context for asking the appropriate questions and for generating viable hypotheses about the conditions that lead to stable, goal-directed behaviors.

The Regulator Paradox: Control and Observation

The challenge that intrigued Wiener and Bigelow as they struggled with the design of systems for guiding artillery during World War II was not the design of simple servomechanisms, as is part of social science lore. Rather, the key problem was to predict the future positions of aircraft. That is, because of the speed and altitude of aircraft, you would not hit an aircraft if you fired at its current position. You had to fire at a point that corresponded to a future position of the aircraft (where the plane would be when the missile arrived). This problem was complicated by the fact that enemy pilots would not fly in simple paths. They would maneuver evasively with the goal to make it difficult to predict where they would be by the time a missile arrived. This problem can best be conceptualized as an observation and prediction

problem. The challenge is to predict the future based on noisy samples of past behavior. This is a problem of distinguishing signal from noise and extrapolating to predict where the aircraft will be when the missile arrives. This is a problem of anticipating or predicting the future. Whereas the goal of the servomechanism (control system) is to minimize “error” (i.e., the deviation between an input goal and behavior output), the goal of the observer is to minimize “surprise” (i.e., the deviation between estimates of target position and velocity based on noisy observations and the actual position and velocity of the target), and the goal of the predictor is to extrapolate that estimate forward in time.

Figure 1.3 illustrates a simple observer system. Note that this is essentially the same dynamical system as in Figure 1.1. The components include a time delay reflecting the time to register the surprise, a gain reflecting the magnitude of adjustment based on the current surprise, and an integrator that effectively sums or concatenates past corrections to determine the next estimation. As in the previous section, let’s assume a fixed time delay that reflects hard constraints on the system. This again leaves the gain as the lone parameter to consider. Let’s consider the input to this system to be a noisy stream of data to which we want to estimate the mean as a reasonable prediction of future samples. For this example let’s consider this a stream of sampled data about the height of people in a particular population as illustrated in Figure 1.4. Further, let’s consider the possibility that in addition to the variability from sample to sample, there is a step change in the mean height of the population (e.g., due to a social change improving diet and access to health care) that occurs during the period of sampling.

In designing an observer for this problem, the goal is to design a system that will not be fooled by the noise (i.e., that will filter out the changes due to noise) but that will be sensitive to real changes in the underlying population (i.e., that will detect

the step change when it occurs). The lone parameter you have to adjust the system is the gain. What value for the gain would be best? Certainly, we know from the discussion in the previous section that if the gain is set too high, the system will become unstable and the estimates will oscillate wildly and will not converge to the input signal. So, we know that there is an upper limit to the gain, but how low should we go? At the setting that is ideal for solving the control problem as it was described in the previous section, the output will follow the input signal very closely. That is, it will respond to every change in input, whether resulting from noise (sampling variability) or from a real change in the underlying distribution, because this distinction was not previously considered in characterizing the input. If the goal is to track the “true signal” based on noisy observations, then determining a satisfactory gain requires simultaneous consideration of both the constraints on control and the constraints on observation. So, it is likely that to reduce the impact of noise, the ideal gain for solving the tracking problem would be lower than if there were no observation noise. In this simplified context, the tracking system in Figure 1.1 has the same structure as the observer in Figure 1.3, so the structure has multiple interpretations.

As in the previous section, there is a fundamental trade-off associated with the gain. Figure 1.4 shows a range of responses for the observer problem as a function of changes in gain. As the gain is lowered, the response to the noise signal will be “smoothed” so that the impact of sampling variability on the estimation will be reduced. However, the price of this “smoothing” is that the observer will be slower to detect real changes when they occur. Thus, in the design of an observer, the engineer must weigh the cost of following the noise against the costs of a slower detection of real changes when they occur. The Kalman filter, which is an optimal solution to the observer




Figure 1.3 This illustrates a simple observer system. The input reflects a noisy signal and the output is an estimate of the signal. In essence, the observer is designed to “filter” out the noise.

problem in a stationary environment. The Kalman filter employs a time-varying gain that minimizes the mean squared error in the estimate.

In essence, this is analogous to a signal detection problem, and the gain parameter will determine the ultimate balance between false alarms (following noise) and hits (detecting real changes). That is, the gain parameter determines the relative sensitivity to signal and noise, as might be captured by the beta parameter in the signal detection model. However, there are important distinctions between the problem as presented here and conventional treatments of signal detection in the social sciences. First, the conventional treatments typically don't consider the factor of time—integration of information from one sample to the next and changes over time in the underlying signal distributions. In essence, conventional treatments treat the signal detection problem as open-loop.

A second important point is that conventional open-loop information processing models and linear closed-loop control models such as the optimal control model (e.g., Pew & Baron, 1978) partition the observation and control processes into separate stages. In a closed-loop model, this elaboration allows different gain settings for the observer and controller. The gain of the observer is chosen to separate signal from noise. The

gain of the controller is chosen to satisfy some effort-accuracy trade-off. This partitioning is advantageous for overall performance. However, it does not capture the true intimacy between perception and action that often exists in closed-loop systems. As Wiener and Bigelow discovered in the process of solving practical control problems like targeting evasive aircraft, there is often an intimate coupling between the observer problem (separating signal and noise to estimate the present position and velocity of an airplane), the prediction problem of extrapolating the path of the aircraft forward in time, and the control problem (aiming projectiles relative to the path of the aircraft). For example, a series of projectiles could be fired either to influence or to constrain the path of the evasive aircraft (i.e., to simplify the observation and prediction problems) as well as to actually hit the aircraft.

In the context of discussing adaptivity, Weinberg and Weinberg (1979) summarize the close coupling between observing and controlling in their Fundamental Regulator Paradox:

The task of a regulator is to eliminate variation, but this variation is the ultimate source of information about the quality of its work. Therefore, the better the job a regulator does, the less information it gets about how to improve.

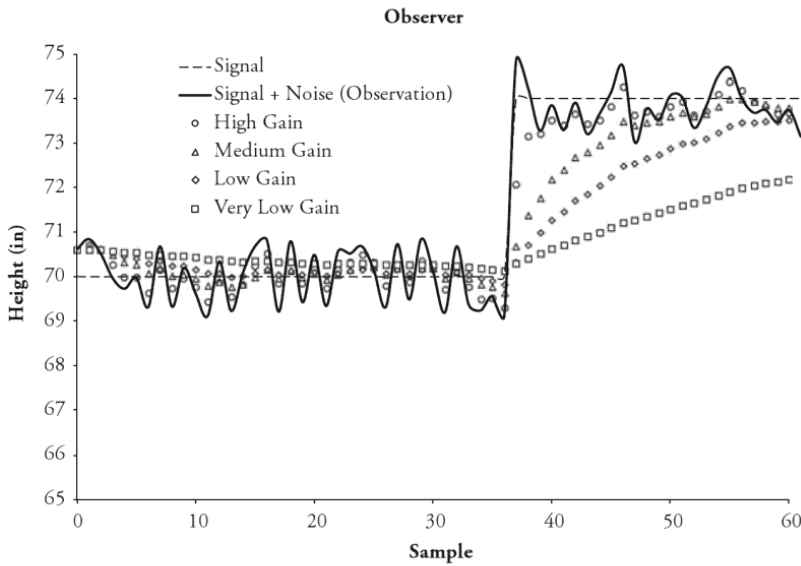


Figure 1.4 This graph illustrates the response of a simple observer as a function of the gain. When the gain is very low, the system will effectively filter out the noise but will be sluggish in responding to changes in the signal. As gain is increased, more of the noise is reflected in the response, but the system is also more responsive to real changes in the signal. At the high gain the output responds to both signal and noise.

or

Better regulation today risks worse regulation tomorrow. (p. 250)

The point is that often cognitive systems and biological control systems must solve the observer, predictor, and control problems simultaneously. That is, they must discover the goals (find the signal in the noise) in real time while simultaneously extrapolating forward in time and managing stable control with respect to those goals. The Weinbergs used the example of controlling a car during the winter to illustrate this coupling. On the one hand, the driver must function as a controller, keeping the car safely on the road. On the other hand, the driver must function as a higher-order observer of system performance—in order to discover potential changes in the control dynamics (e.g., resulting from changes of the road surface such as black ice). In balancing the demands of these two tasks, good drivers learn to act to both steer the car and to test for changes in road conditions (e.g., put in test signals such as jiggling the steering wheel or tapping the brakes). Note that the test signals produce “error” with regards to the control problem, but they provide information with regards to the changing dynamic context (e.g., detecting control-relevant changes in the situation). This information may be critical for maintaining stable control when road conditions change. Gibson (1979/1986) used the terms “performatory action” and “exploratory action” to distinguish between those actions motivated by the demands of control and those actions motivated by the demands for information.

To reiterate, in designing automatic control systems, engineers typically “solve” the control problem and the observer problem offline and then implement the solutions as an automatic control system (with gains tuned to reflect both the signal-to-noise and the stability constraints of the problem). This works well if the environments are relatively stable (i.e., stationary). However, cognitive and biological systems have to design themselves while simultaneously interacting with the problem situation. For example, the new driver or new pilot must discover the appropriate gains with respect to 1) signal-and-noise (i.e., determine what aspects of variability to filter out) and 2) good control (i.e., efficient, stable correction of errors) while engaged in the control task. These systems must learn by doing. These systems must self-organize. Further, to survive, the cognitive and biological systems must be prepared to adapt to changing contingencies.

That is, they cannot assume a stable environment. For example, the operators of the anti-aircraft artillery cannot assume that the enemy pilots will not change their evasive tactics, and the operators may even act to influence those tactics.

Today, control theory and information theory are typically treated as distinct fields of research. This is in part due to the different analytic tools (differential equations for control theory; probability statistics for information theory). However, both fields emerged as a result of the work of Wiener and Bigelow on the problem of predicting aircraft movements in order to target them. Note that the subtitle of Wiener’s (1948/1968) classical work *Cybernetics* was *Control and Communication in the Animal and the Machine*. In closed-loop systems, observation and control are two sides of a single coin!

The main points for this section were:

1. To show the feedback dynamic from the perspective of information processing or perception. In this case, the gain functions to tune the filtering properties of the dynamic—in order to distinguish signal from noise based on observations over time.
2. To suggest that in most natural systems the control and observer problems are intimately linked. There could be a single gain parameter to meet both the constraints on observation and control in a simplified context, or there could be two or more independent gain parameters for observer and controller that together determine the overall responsiveness and stability of the closed-loop system. In addition, the actions of the system may have multiple goals of influencing the observation/prediction process and well as tracking the input signal.
3. Finally, it is important to appreciate that research into closed-loop systems must be guided by the intuitions of both control and information theory.

The Comparator Problem

In a control or observer system, the comparator is the point at which the output is fed back and “compared” with the input to compute an error (or surprise) signal. In typical engineered control systems such as those illustrated in Figures 1.1 and 1.3, the system is designed so that the comparator simply involves subtraction of one signal from another to get a third signal. That is, all signals are in a comparable currency that allows subtraction of one from another to get the third. However, consider




Figure 1.5 A second-order control system in which the gain determines the acceleration of the response (output).

the novice pilot learning to land. The goal might be a particular approach path or simply a soft contact at a particular region of the runway. The feedback would be in the form of optical flow through the windows and/or information presented on the cockpit instruments. These two very different types of signals must somehow be compared in order to specify appropriate movements of the controls (stick, throttle, rudders).

It should be clear that for the pilot, and more generally for most cognitive or biological control systems, the signals involved in the comparator process may be in diverse forms or currencies. Thus, comparing feedback to intentions in order to specify actions is not a trivial process. In fact, this is a central issue for control theory—to determine the dimensionality of the state space or, in other words, to identify what variables must be fed back in order to guide action in a particular situation (e.g., as a function of different vehicle dynamics). This is also probably the central issue for skill development—attuning to the feedback that specifies the appropriate actions with respect to the opportunities and consequences (e.g., E. J. Gibson, 1969; Ericsson & Charness, 1994). In Gibsonian (1979/1986) terms, this is the problem of specification of affordances.

Figure 1.5 shows a simple feedback system to illustrate the comparator problem. This system has a gain and two integrators in the forward loop. As with the other systems, the gain determines the sensitivity to error. Because of the two integrations in the forward loop, the output from the gain element determines

the acceleration of the output. This is a dynamic that is consistent with most movement tasks (e.g., vehicle control or body movement) in a world governed by inertia. For example, the initial response of deflection of the accelerator or brake is an initial change in the acceleration or deceleration for your car.

What do you suppose the response of this system would be to a step input, and how might this response change as a function of the lone parameter in the forward loop—the gain? Is there any gain value that will result in an asymptotic approach to the step target? Somewhat surprisingly in the context of naive discussions of feedback systems, the answer is “No.” Here again is a situation that illustrates that feedback is not sufficient to insure convergence with the input. In fact, the response of this system to a step input is a sine wave output. The speed of oscillation of this sine wave (i.e., its frequency) is determined by the value of the gain parameter. A higher gain produces a higher-frequency response. There is NO value of gain that will lead to convergence of output with the input target!

Figure 1.6 shows an alternative system that includes feedback of both the output position and the output velocity. This system will result in an output that will converge to the input target, if the feedback of position and velocity are combined with the appropriate weights. Heavy relative weight on the velocity component will lead to “sluggish” or conservative approaches to the target. Less relative weight on velocity will lead to more aggressive approaches—with a damped oscillation at very low

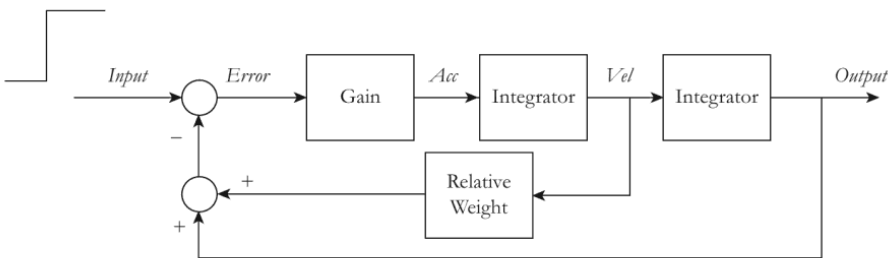


Figure 1.6 Stable goal following can be achieved in a second-order system if feedback includes both position and velocity of the output.

weights and, as described above, a non-converging oscillation at zero weight on velocity. Depending on the relative weights on position and velocity, this system will achieve a similar range of responses (from sluggish to oscillatory) as illustrated in Figure 1.2.

A general implication of the behavior of the systems illustrated in Figures 1.5 and 1.6 is that for control of an inertial system (e.g., a car), position feedback alone is insufficient. The system must have feedback about both position and velocity in order to control the vehicle. For example, in order to stop the car in front of an obstacle on the road (e.g., stopped line of traffic), a driver must take into account both distance to the obstacle and the speed of approach. Current research suggests that for visual control of locomotion, this information is specified in terms of the angular extent (e.g., visual angle of the taillights of the preceding vehicle) and the angular velocity (e.g., rate of expansion of the taillights) of the object in the visual flow field (e.g., Lee, 1976; Smith, Flach, Dittman, & Stanard, 2001; Flach, Smith, Stanard, & Dittman, 2004).

In most natural situations, cognitive systems must deal with many different potentially useful sources of feedback. Figure 1.7 provides a simplified illustration of the multiple variables that need to be considered in controlling the lateral position of an aircraft when trying to track a specific

approach path. The top portion of Figure 1.7 illustrates the multiple “state variables” that a controller must be aware of in order to control effectively. The input from the pilot directly affects the position of the aileron, which is integrated (first-order lag) to determine the roll rate, which in turn is integrated to determine the roll angle, which is integrated to determine the heading angle or turn rate, which ultimately is integrated to determine the lateral position of the aircraft. The output of each integration represents a “state variable” of the aircraft that must be fed back in order to achieve reliable control. Note that the lateral control problem is only a subset of the variables that need to be considered in landing.

The lower portion of Figure 1.7 illustrates how each of these state variables might be fed back and combined to determine the appropriate control adjustments. Each loop has a specific gain that in effect reflects the relative weighting of each of the state variables in determining the next adjustment of aileron position. In designing an effective control system, the engineer would set the gains in each loop to reflect the dynamics of a particular aircraft. Similarly, a pilot who is learning to fly the aircraft must discover the appropriate mappings from the perceptual information associated with each variable (e.g., properties of optical flow or instrument

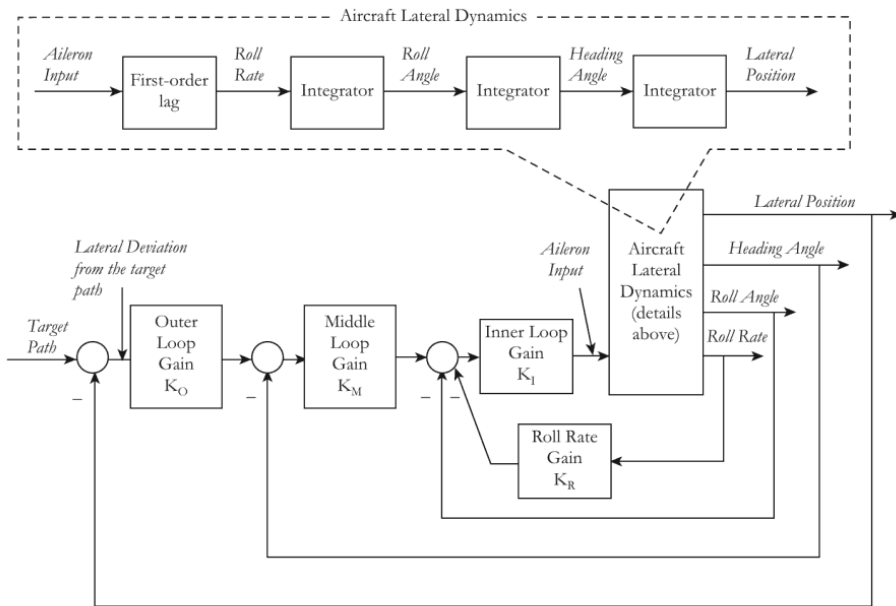


Figure 1.7 The top diagram provides a highly simplified representation of the aircraft lateral dynamics and the state variables involved in controlling lateral position of an aircraft (e.g., in trying to track the target approach path to an airport). The bottom diagram illustrates the multivariable control problem. Each of the state variables are perceived, weighted, and combined to determine the appropriate control actions.

readings) and the adjustments of his or her control interface (e.g., control stick).

In addition to considering the multiple state variables associated with the lateral and vertical positions of the aircraft, pilots must consider potentially conflicting goals (e.g., aborting the approach when a taxiing aircraft encroaches on the runway), and many potential actions or means for moving toward those goals (e.g., whether to make stick or throttle adjustments to correct for excessive air speed). Thus, the comparator problem involves attuning to many potentially useful sources of feedback; it involves setting priorities to appropriately trade off the potential consequences of multiple competing goals; and it involves choosing among many potential means for reducing the error between output and intention. These multiple degrees of freedom in terms of multiple consequences, multiple sources of information, and multiple potential control actions emphasize the intimacy between observation and control introduced in the preceding section.

Also, it is important to understand that the setting of the gains in Figure 1.7 that result in satisfactory control depend on the particular aircraft dynamic. However, the aircraft dynamic itself may change as a function of context. For example, the dynamics can change as a function of both speed and altitude. For example, gains that lead to stable control at low altitudes may lead to instabilities at higher altitudes. In the design of automatic control systems, engineers address this by designing adaptive control systems. As illustrated in Figure 1.8, an adaptive control system is capable of adjusting the gains on the inner loop dynamics as a result of monitoring the context or the quality of inner loop responses.

In Figure 1.8, the thin (wire) arrows represent the flow of signals or information that is then acted on or processed according to the workings (i.e., the transfer function—e.g., control gains) of each box. However, the fat arrows that close the outer loops are signals that change the properties of the boxes (e.g., change the control gains or the expectations within the boxes). Note that the inner loop includes three modes of action: performatory actions are intended to reduce error, exploratory actions are intended to test hypotheses, and anticipative actions reflect direct action to achieve a goal. The anticipative path reflects direct response to the reference (i.e., not dependent on error feedback). This is an open-loop path from the reference or goal and could reflect actions that are shaped by previous experience with the plant dynamics (e.g., what has often been

called a mental model or schema). That is, these are responses in anticipation of a consequence, rather than in response to error feedback. Actions may fulfill any or all of these roles at any moment.

Figure 1.8 illustrates three different ways that an engineer might close the outer loop to achieve stable adaptive control. First, the adaptation might be a direct function of the changing context, as illustrated in the outermost loop. For example, the engineer might compute the appropriate gains for different altitudes and preprogram these different gains into the automated control system. The gains would be changed as a function of a direct measure of the appropriate context variable (e.g., the altitude). This is typically called “gain scheduling.” In human performance, this path may be representative of the phenomenon of context sensitivity. That is, the strategy for controlling action or the expectations of the human agent may change as a function of the situation. For example, data suggests that a teenage driver with a parent in the car is one of the safest drivers, while a teenage driver in a car with other teenagers is one of the most dangerous drivers.

The next outer loop (hypothesis) in Figure 1.8 reflects conscious exploration of the dynamics through exploratory actions. In engineered systems this might involve a low-amplitude test signal that is constantly input to the plant (i.e., dithering). This input is designed to have minimum consequences relative to the performance objective (i.e., to produce minimal error). However, the changes in the properties of the output from this signal can be information relevant to detecting changes in the plant dynamics. Deviation in the output related to the dithering can be fed back and used to adapt the control gains (and the expectations). Skilled human drivers use a similar strategy to test for possible changes in the driving dynamics due to changing road conditions. They may dither the steering wheel to test for changing traction (Weinberg & Weinberg, 1979). The reference signals for this loop are labeled “local expectations” to reflect that this loop reflects explicit tests of local hypotheses. In this loop the human agent is acting as a test signal generator and observer—in order to detect control-relevant changes.

The innermost of the outer loops (surprise) represents an approach to adaptive control that engineers call “model reference” control. With this style of adaptive control, a normative model of the plant dynamics can be simulated in parallel with the actual performance. For example, this model might be a simulation of aircraft performance at the typical

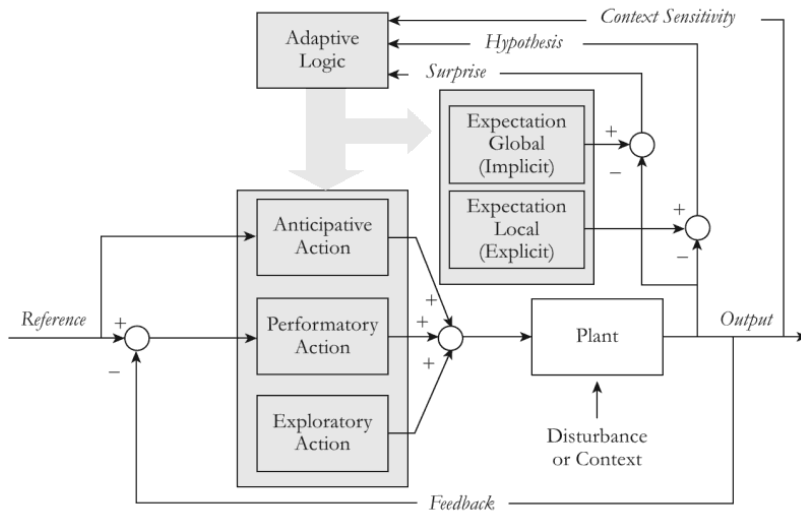


Figure 1.8 This diagram combines an inner (control) loop with outer (adaptive or metacognitive) loops to illustrate three styles that engineers use in the design of adaptive control systems.

altitude. Expectations based on the normative simulation can then be compared with the actual behavior of the vehicle, and deviations from expectation can be fed back to adapt the control strategy (and the expectations). Again, the engineer’s “simulation” may be somewhat analogous to what cognitive scientists refer to as knowledge or a mental model. This “model” reflects integrated experiences from the past that provide a backdrop that experts can use to assess situations. In many cases, this “internal model” operates implicitly—so that experts may not become aware of the expectations until there is a mismatch. And even when the mismatch is noticed, the feedback may be experienced only as a vague sense that something is not normal. Thus, this might be the basis for intuitive aspects of expertise (Klein, 2004).

Overall, the outer loops in Figure 1.8 may provide a constructive way to think about the general phenomena of “metacognition,” that is, self-awareness or our ability to monitor and critique our own performance. This is another layer of closed-loop, iterative processing in which the processes are simultaneously shaping their responses and being shaped by those responses. Again, this reflects the self-organizing aspect of cognitive systems. These are systems that are capable of learning from their mistakes.

Adaptive control systems where outer loops change parameters of inner loops are inherently nonlinear. The stability of these systems becomes a much more difficult analytic problem for control theory, particularly considering that all three outer loops may be acting simultaneously. These systems can be trapped in local minima (e.g., converge on a

degenerate model of the plant; superstitious behavior), and they are vulnerable to butterfly effects (e.g., a small change can cascade, having dramatic effects with regards to stability).

Finally, it should be apparent that as we increase the number of variables that must be fed back and as we add outer loops that change inner loops, the complexity of the computations involved expands rapidly, thus raising questions with respect to the capacity of the human agent to manage this complexity. This issue will be addressed in the next section.

The main points of this section are:

1. To increase appreciation of the natural complexity of the comparator process whereby feedback is compared with intentions in order to specify corrective actions. The comparator process is typically represented as simple subtraction. This formalism greatly trivializes the problem faced by cognitive systems.
2. To reinforce the intimacy of observation and control in closed-loop systems. The closed-loop dynamic demands that questions of intention, perception, and action be framed in relation to each other. A program that studies these as independent stages in a linear causal stream will result in trivializing the dynamics of cognition.
3. To reinforce the fact that despite the common misconception, feedback does not insure stable convergence of output to the target input due to the complexity of multi-loop dynamics and nonstationary (i.e., changing over time) parameters.

Smart Mechanisms and Situation Awareness

There are two important consequences of having multiple degrees of freedom in terms of intentions, feedback information, and possible actions. On the positive side, multiple degrees of freedom mean that there are many ways to “skin a cat.” That is, there will typically be multiple ways to satisfy an intention. This reflects a range of goal trade-offs that might be acceptable, multiple potentially redundant sources of information relative to guiding action, and multiple combinations of actions that move the system toward satisfactory goals. On the negative side, multiple degrees of freedom can greatly increase the computational demands on the system. That is, the computational load associated with the comparator problem will be a rapidly increasing function of the number of variables involved—due to the potential for multilevel interactions among these variables and across inner and outer control loops.

It is clear that the human is a limited bandwidth information processor. That is, there is a limit to how many independent things the system can attend to or be aware of at the same time (e.g., Broadbent, 1971; Miller, 1956). Thus, there will be limits to the complexity of the comparator problem that humans (and other biological systems) are able to solve. In general, for “good” or “smart” control, the system should prefer solutions that both satisfy the intentions and minimize the computational demands (i.e., demands on awareness).

Bernstein (1967) was one of the first to draw attention to the awareness or computational constraints on feedback control in the context of motor skills. He noted that from the many possible means for achieving a goal, skilled athletes tended to choose solutions that reduced the degrees of freedom that needed to be monitored (or controlled) in real time. This was typically accomplished by locking out or constraining other potential degrees of freedom. This is easily illustrated by considering the swing of a skilled golfer when driving a golf ball. The way that the golfer holds the club, addresses the golf ball, and moves during the swing is explicitly designed to lock out many degrees of freedom (e.g., keeping head fixed with eyes on the ball and keeping the elbow of the left arm straight for a right-handed golfer). The result of locking out or constraining many potential degrees of freedom is that the complexity of the comparator problem is greatly reduced. It reduces the number of variables that need to be monitored/controlled in real time (i.e., attended or adjusted) in order to control convergence to a satisfactory solution. For example, if the elbow position of the left

arm were also changing during a swing, its interactions with all the other variables (e.g., wrist angle) would have to be monitored and compensated for in order to reliably make contact with the ball. The term “coordinative structures” has been suggested for the patterns of constraint seen in skilled motor activities (Turvey, 1977).

Runeson (1977) contrasted the solutions such as coordinative structures with some of the early conventional approaches to engineering and general problem-solving processes. The conventional approaches tended to start with a fixed general coordinate system (e.g., orthogonal three-dimensional spatial coordinates), and all control problems were then organized with respect to those dimensions (e.g., describing the motion of all arm components with respect to the positions and velocities in this orthogonal space). The result of using this fixed coordinate space is that many common natural motions required relatively complicated descriptions; thus, very complex computations were implied. To simplify computations, engineers typically choose solutions that are more compatible with the fixed coordinate system. This results in very unnatural, “robotic” motion. Runeson called the systems resulting from reliance on a single, fixed coordinate system “rote mechanisms,” suggesting that these reflected brute, one-size-fits-all solutions.

Alternatively, Runeson suggested that biological systems change the coordinate system to reflect the intrinsic constraints on the degrees of freedom (i.e., the coordinative structures). In these cases, a coordinate system was chosen to make the description of the solutions simpler—in effect reducing the computational demands. Runeson called systems that organize the comparator processes around “intrinsic” rather than “extrinsic” coordinates “smart mechanisms.”

The constructs of “coordinative structure” and “smart mechanism” have important implications for understanding the dynamics of cognitive systems. Much of the research on perception defines “space” with respect to a fixed coordinate system, where human performance is evaluated with respect to fixed rulers (e.g., length measured in cm) and fixed coordinate systems (e.g., height, length, depth perception). This classic approach is designed around open-loop tasks where a stimulus is presented (e.g., an object of a particular size and position in the world), and subjects are asked to make passive judgments (e.g., how big is the object? or how far away is it?). Great care is typically taken to insure that the subjects’ responses do not change the

stimulus (e.g., stabilizing the eyes using chin rests or bite bars).

The constructs of coordinative structure and smart mechanism suggest that the closed-loop system may self-organize around the intrinsic geometry of the problem (e.g., vehicle dynamics). This suggests that intrinsic aspects of the situation (e.g., imminence of collision relative to the braking dynamics) may be more “basic” to perception than the extrinsic coordinates used by physicists to describe space (e.g., depth perception). In other words, this suggests that “space” in the classical extrinsic sense is not the dominant coordinate frame for perception. That is, the perceptual system is not designed to “see” space, but rather to detect those properties most directly relevant to guiding action (e.g., the imminence of collision). It is designed (or has evolved) to solve the comparator problem. Gibson (1979/1986) has argued for this hypothesis in terms of his construct of “direct perception of affordances.” That is, a smart mechanism organizes itself so that the functionally relevant properties (i.e., the affordances) are specified in a simple or direct way to reduce the computational demands. This idea, which once seemed quite radical, becomes quite plausible when considered in the context of feedback control.

In the spirit of Gibson’s approach, the challenge is not to understand “space perception” per se. Rather, the challenge is to understand the ability for experts to discover those coordinates that provide the most direct mapping across the dimensions of intention, feedback, and action. In other words, the challenge is to find the smartest solutions to the comparator problem (e.g., the simplest solution that satisfies the demands of the situation). While in some cases the solutions might be based on global constraints (suggested by the search for optical invariants), recent research suggests that even in relatively simple perceptual judgments, people are likely to utilize more local constraints of specific task environments when they are locally reliable (e.g., kinematic correlates to the specifying mass ratios for predicting the results of collisions) (Jacobs, Runeson, & Michaels, 2001). This suggests that even solutions to basic perception tasks may be more heuristic in nature than had previously been thought by those following a Gibsonian approach. That is, people will often organize behavior around local constraints, rather than always utilizing global invariants. However, Gibson’s main point remains—the key to understanding the coupling between perception and action is to discover the intrinsic coordinates that allow direct mapping between intentions, consequences, and action.

The key point of this section is to reframe the challenge of cognition. In classical approaches the challenge of cognition is framed as the ability to construct a valid internal model of the world (that then becomes the basis for motor control and decision making). This “internal model” is typically judged relative to extrinsic physical models of space (e.g., standard rulers and coordinate systems). We suggest that the problem of cognition is simply to guide or direct successful action. That is, the function of cognition is to solve the comparator problem. This requires that our understanding of “awareness” (e.g., internal models) be grounded in an understanding of “situations” (e.g., Flach & Rasmussen, 2000; Flach, Mulder, & van Paassen, 2004). The “test” of a belief is not based on classical induction or deduction, but rather it is based on the consequences of actions based on that belief.

It is important to understand that this mapping between perception and action is not a trivial problem (e.g., consistently hitting a nail is not as trivial as suggested by Miller et al.’s TOTE example). It is very likely that as a result of searching for and discovering smart control mechanisms for the comparator problem, we build internal representations of the world that might generalize to more abstract models that guide scientific exploration and discovery and that create new ways for interaction. However, from the circular systems perspective, the abstract models of science are a product of the coupling of perception and action, not a necessary prerequisite for the coupling. We discover the world through acting in it.

Generalizing to Problem Solving and Decision Making

The relevance of the dynamics of closed-loop systems is easy to discover in the context of programs of research on perceptual-motor coordination (e.g., manual control). In that research context, there is an obvious mapping between the normative models (e.g., optimal control models and automated control systems) and the cognitive phenomena (e.g., piloting an aircraft), and there is a history of work utilizing the analytic tools of control theory (e.g., frequency domain analysis) to describe human performance (e.g., Sheridan & Ferrell, 1974; Jagacinski & Flach, 2003). However, do these principles have implications more generally for cognitive systems? Is the logic of circles also relevant to phenomena associated with decision making and problem solving?

We believe that the answer to this is obviously “Yes.” This belief is supported by observations of

naturalistic decision making that discovered that in many decision contexts there is an intimate link between “recognition” (i.e., perception) and “choice” (i.e., action) (Klein, 1993). For example, choices (decisions/actions) of a fire ground commander (e.g., where to direct the hoses) are primed by recognition processes or feedback (e.g., observations of the fire relative to the goal priorities—saving lives and property). It is also supported by work on problem solving in the wild (Hutchins, 1985). For example, Hutchins (1985) describes how the solutions to navigation problems are shaped by the available tools and representations (e.g., maps). Again, we believe that these analyses of “cognition in the wild” are unpacking the comparator problem. That is, these analyses are discussing the relations between input (e.g., intentions/goals), output (e.g., consequences, feedback), and action. Different strategies can then be evaluated relative to whether they lead to satisfaction of the goals (i.e., whether the system converges to stable solutions) and relative to the computational load (i.e., how smart are the computational mechanisms?). Finally, a big part of the story is how experts in natural work domains leverage intrinsic problem constraints to improve efficiency (i.e., coordinative structures).

Cognitive theory and research continues to be dominated by the logic of open-loop causal reasoning, despite the growing evidence that cognitive phenomena in nature are typically closed-loop. For example, human decision making continues to be evaluated against normative models that are based on open-loop logic (induction, deduction). However, the logic of abduction that was first proposed by Charles Sanders Peirce (1978) provides an alternative basis for rationality that is more consistent with the logic of circles. For Peirce, the test of a hypothesis was not the form of the argument (as suggested by classical logic), but rather, the test of the hypothesis was the pragmatic consequences of actions directed by that hypothesis.

The abductive process described by Peirce is directly analogous to the observer illustrated in Figure 1.3. That is, an abductive system is a system designed to eliminate surprise. In an abductive system, as in an observer, beliefs that lead to accurate perceptions and predictions about the world (i.e., converge in ways that reduce surprise) are strengthened. However, remember that this observation process is not independent from the process of control. Thus, the real test of our beliefs in an abductive system is whether the actions guided by these beliefs lead to satisfactory consequences (i.e., reduce error

with respect to intentions). In essence, the targeting problem that challenged Wiener and Bigelow was a problem of abduction. The problem was to anticipate the actions of an aircraft in order to successfully target it. The ultimate test of any belief/guess about the behavior of the aircraft was the consequences of an action (either firing projectiles or simply comparing predicted and actual behavior) that were then fed back, compared, and integrated to both change the belief (prediction) and guide the next prediction and/or action.

Conclusion

The primary motive for this chapter is that we believe that cognition is a closed-loop phenomenon. Thus, the logic of circles is fundamental. We believe that many of the conundrums in cognitive science and the failures of that science to inform design are the result of framing the questions in the context of simple open-loop causal systems. Although Wiener’s cybernetic hypothesis has helped to inspire the cognitive revolution, and although feedback loops are often included in the images we create to represent cognitive systems, we believe that for the most part the logic of circular systems has not been fully appreciated or applied. For the most part, cognitive theory and research continues to be framed in an open-loop context.

This is not simply a problem for cognitive science, but it seems to be integral to the more general logic of experimental science, where experiments are explicitly designed to explicate open-loop cause-effect relations. Note that in considering a circle as a whole, there is no unambiguous direction for distinguishing cause from effect. In a recursive circular dynamic, what we see determines what we do, while simultaneously what we do determines what we see. Observation and control are intimately coupled. There is no basis for one (either perception or action) to have causal precedence over the other.

There is a tendency for experimental science to organize around general, extrinsic coordinate systems that are at least implicitly accepted as ground truth. These truths become the ruler against which behavior is measured. For example, perception of size is judged against a standard measure in centimeters. Or choices are judged against the prescriptions of inductive or deductive logic. However, circular systems are self-organizing. This means that understanding will generally depend on the ability to discover the intrinsic coordinates (or standards) that are structuring that organization. Thus, the question is not “how big?” or “how far?” in any absolute

sense (e.g., X cm), but rather whether it is graspable, or pass-through-able, or whether collision is imminent relative to the capacity to brake. In a circular system, the value of any variable (intention, action, or perception) must be “measured” relative to the other variables. The meaning of “too close” is contingent on the action dynamics (i.e., the capacity for braking or evasive maneuvering). The question is not whether the reasoning is “sound” in an absolute sense (e.g., valid relative to normative prescriptions), but rather does the reasoning lead to successful action?

In the classical approach, performance is gauged with respect to ideal norms derived from abstract mathematical models that were intended to generalize across a myriad of different situations. We believe such abstract norms can be very important for bounding the space of possible behaviors, but that they are often not the appropriate gauges for understanding the dynamics of actual behavior. Actual behavior is grounded in the pragmatics of situations, and while an artful application of mathematics may be essential for describing the resulting patterns and constraints, it needs to be particularized to distinct situations. Actual behavior of circular systems is situated. Closed-loop systems are adapting to their environments, and they are simultaneously adapting (i.e., changing) their environments. This coupling between environment and organism is reflected in the dynamics of prey-predator systems, where the size of each population is shaping and being shaped by the size of the other population. This coupling and the resulting self-organization are also nicely illustrated by Kugler and Turvey’s (1987) metaphor of insect nest building. In this system, insect behavior creates the pheromone landscape at the same time the pheromone landscape is shaping insect behavior. The resulting structure of the insect nest is a product of this coupling.

In this chapter, our goal was to use a few very simple examples to help people to appreciate the dynamics of closed-loop systems. These examples were chosen to highlight differences between the expectations developed from an open-loop theory of behavior and to increase appreciation for the complexity of closed-loop systems. We intentionally chose very simple closed-loop systems so that the examples would be understandable and to minimize any need for complex analysis. However, it is important to realize that we are not suggesting these examples as specific “models” or “metaphors” of cognition, although we do believe that intuitions developed through understanding these simple

circular dynamics may be important in shaping a theory of cognition that captures the dynamics of particular situations.

We believe that the evidence for circularity in natural cognitive phenomena is pervasive. Although one can find local behaviors that are open-loop (e.g., ballistic braking), even in these cases the local actions will typically be components of a larger circular dynamic (e.g., safe driving). However, the point here is not to fully elucidate cognition, but to elucidate the logic of circular systems. Our hope is that a deeper appreciation and understanding of closed-loop systems will enrich theory, experimental practice, and application of cognitive science. We hope that a deeper appreciation of closed-loop systems will lead cognitive scientists to ask better questions.

Finally, we believe that posing a question correctly will take us much further toward finding satisfactory answers to that question. We are very optimistic that trends in cognitive science (e.g., associated with neural nets, dynamical systems, artificial life, etc.) and in cognitive engineering (e.g., associated with ecological interfaces and semantic computing) suggest that appreciation for the circularity inherent in cognitive phenomena is growing. However, unfortunately, this appreciation is hampered by pervasive temptations to trivialize the dynamics of circles in order to satisfy conventional assumptions about the open-loop nature of explanation and the open-loop nature of experimental inference. Although most programs in cognitive science require students to take experimental methods courses (e.g., to learn analysis of variance and regression), we are aware of no academic program in cognitive science that requires a course to introduce the dynamics of closed-loop systems.

Future Directions

We suggest three challenges for the future associated with theory, methodology, and practice. On the theoretical side, it is important to move beyond trivial control metaphors to more carefully consider the dynamics of closed-loop systems. There are two levels of theory to consider. At the meta-level, we need to give up the domino theory of open-loop causality to consider general models of complex, dynamical systems. At the base theoretical level, we need to attend to the situated (e.g., Suchman, 1987) or embodied (e.g., Clark, 1997) dynamics of cognition. We need to move beyond an exclusive focus on formal logic and other context-free models of rationality to consider the pragmatic, ecological

constraints of perception and action (e.g., abduction) (e.g., Todd & Gigerenzer, 2003).

Methodologically, it is important to broaden the empirical base. This includes attending more to naturalistic observations of cognition in complex work domains (e.g., Hutchins, 1985). It also means considering more representative designs for controlled experimentation (e.g., Kirlik, 2006). The representative designs need to provide more degrees of freedom to the participants with respect to both the means and ends for performance. If the laboratory tasks are designed around the simple servomechanism metaphor, then of course the humans will adapt and behave like a simple servomechanism. If we want to explore the creative, adaptive capacity of human cognition, then we have to create laboratory situations that invite creativity and adaptation. Synthetic task environments provide a unique opportunity here (Flach, Schwartz, Courtice Behymer, & Shebilske, 2010). In these environments it is possible to empirically link performance at the micro-level (e.g., reaction time to specific events) to functional ends at a macro-level (e.g., success in domain terms, such as successfully completing a mission). The capacity to empirically link variations at the micro-performance level with more global functional consequences will be critical to modeling self-organizing dynamics.

Finally, at a practical level, it becomes necessary for enhanced collaborations between those who focus on awareness (e.g., cognitive scientists) and those who focus on situations (e.g., engineers, domain experts). Without cooperation and mutual respect across disciplines, it will never be possible to achieve a deep understanding of situation awareness. This will require that we escape from the view where basic and applied sciences are seen as competitors in a zero sum game. We need to embrace the spirit of Pasteur's quadrant, where the goals of theory and practical utility are respected and valued equally as complementary components of a mature science (Stokes, 1997).

References

- Bernstein, N. A. (1967). *The control and regulation of movements*. London, England: Pergamon Press.
- Broadbent, D. A. (1971). *Decision and stress*. New York, NY: Academic Press.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.
- Conway, F., & Siegelman, J. (2005). *Dark hero of the information age: In search of Norbert Wiener, the father of cybernetics*. New York, NY: Basic Books.
- Ericsson, K. A., & Charness, N. (1994). Expert performance: Its structure and acquisition. *American Psychologist*, 49(8), 725–747.
- Flach, J., Mulder, M., & van Paassen, M. M. (2004). The concept of the “situation” in psychology. In S. Banbury and S. Tremblay (Eds.), *A cognitive approach to situation awareness: Theory, measurement, and application* (pp. 42–60). Aldershot, England: Ashgate.
- Flach, J. M., & Rasmussen, J. (2000). Cognitive engineering: Designing for situation awareness. In N. Sarter & R. Amalberti (Eds.), *Cognitive engineering in the aviation domain* (pp. 153–179). Mahwah, NJ: Erlbaum.
- Flach, J. M., Schwartz, D., Courtice A., Behymer, K., & Shebilske, W. (2010). Synthetic task environments: Measuring macrocognition. In E. Patterson & J. Miller (Eds.), *Macrocognition metrics and scenarios: Design and evaluation for real-world teams* (pp. 201–284). Aldershot, England: Ashgate.
- Flach, J. M., Smith, M. R. H., Stanard, T., & Dittman, S. M. (2004). Collisions: Getting them under control. In H. Hecht & G. J. P. Savelsbergh (Eds.), *Time to contact, Advances in Psychology* (Vol. 135, pp. 67–91). North-Holland, Amsterdam: Elsevier.
- Gibson, E. J. (1969). *Principles of perceptual learning and development*. New York, NY: Appleton-Croft.
- Gibson, J. J. (1979/1986). *The ecological approach to visual perception*. Hillsdale, NJ: Erlbaum.
- Hutchins, E. (1985). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Jacobs, D. M., Runeson, S., & Michaels, C. F. (2001). Learning to visually perceive the relative mass of colliding balls in globally and locally constrained task ecologies. *Journal of Experimental Psychology: Human Perception and Performance*, 27(5), 1019–1038.
- Jagacinski, R. J. (1977). A qualitative look at control theory as a style of describing behavior. *Human Factors*, 19, 331–347.
- Jagacinski, R. J., & Flach, J. M. (2003). *Control theory for humans: Quantitative approaches to modeling performance*. Mahwah, NJ: Erlbaum.
- Kirlik, A. (2006). *Adaptive perspectives on human-technology interaction*. New York, NY: Oxford University Press.
- Klein, G. A. (1993). A recognition-primed decision (RPD) model of rapid decision making. In G. A. Klein, J. Orasanu, R. Calderwood, & C. E. Zsombok (Eds.), *Decision making in action: Models and methods* (pp. 138–147). Norwood, CT: Ablex.
- Klein, G. A. (2004). *The power of intuition*. New York, NY: Random House.
- Kleinman, D. L., Baron, S., & Levison, W. H. (1971). A control theoretic approach to manned-vehicle systems analysis. *IEEE Transaction on Automatic Control*, AC-16, 824–832.
- Kugler, P. N., & Turvey, M. T. (1987). *Information, natural law, and the self-assembly of rhythmic movement*. Hillsdale, NJ: Erlbaum.
- Lee, D. N. (1976). A theory of visual control of braking based on information about time-to-collision. *Perception*, 5, 437–459.
- McRuer, D. T., & Jex, H. R. (1967). A review of quasi-linear pilot models. *IEEE Transactions on Human Factors in Electronics*, HFE-8, 231–249.
- Miller, G. A. (1956). The magical number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 81–97.
- Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior*. New York, NY: Henry Holt and Company.

- Peirce, C. S. (1978, August). Deduction, induction, and hypothesis. *Popular Science Monthly*, 13, 470–482.
- Pew, R. W., & Baron, S. (1978). The components of an information processing theory of skilled performance based on the optimal control perspective. In G. E. Stelmach (Ed.), *Information processing in motor control and learning* (pp. 71–78). New York, NY: Academic Press.
- Rosenblueth, A., Wiener, N., & Bigelow, J. (1943). Behavior, purpose, and teleology. *Philosophy of Science*, 10, 18–24.
- Runeson, S. (1977). On the possibility of “smart” perceptual mechanisms. *Scandinavian Journal of Psychology*, 18(1), 172–179.
- Sheridan, T. B., & Ferrell, W. R. (1974). *Man-machine systems: Information, control, and decision models of human performance*. Cambridge, MA: MIT Press.
- Smith, M. R. H., Flach, J. M., Dittman, S. M., & Stanard, T. W. (2001). Alternative optical bases for controlling collisions. *Journal of Experimental Psychology: Human Perception and Performance*, 27, 295–410.
- Stokes, D. E. (1997). *Pasteur’s quadrant: Basic science and technological innovation*. Washington, DC: Brookings Institutions Press.
- Suchman, L. A. (1987). *Plans and situated action*. Cambridge, England: Cambridge University Press.
- Todd, P., & Gigerenzer, G. (2003). Bounding rationality to the world. *Journal of Economic Psychology*, 24, 143–165.
- Turvey, M. T. (1977). Preliminaries to a theory of action with reference to vision. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting and knowing: Toward an ecological psychology* (pp. 211–265). Hillsdale, NJ: Erlbaum.
- Weinberg, G. M., & Weinberg, D. (1979). *On the design of stable systems*. New York, NY: Wiley.
- Wiener, N. (1948/1961). *Cybernetics: Or control and communication in the animal and the machine*. Cambridge, MA: MIT Press.

Christopher D. Wickens

Abstract

This chapter describes attention in cognitive engineering and in design by metaphors of **the filter**, that selects incoming sensory information for perception, and the **fuel** that supports all stages of information processing with a limited supply of resources, and that therefore limits multi-tasking. We describe applications of the filter to noticing events, alarm design, supervisory control and display layout, display integration, and visual search. We then consider two aspects of multi-task performance: when fuel is available to support concurrent processing, predicted by a multiple resource model, and when the task demands are sufficiently high, as to force sequential processing, and consideration of task and interruption management strategies. Finally, we consider the role of mental workload in automation and situation awareness. Where relevant, the chapter highlights the role of computational models.

Key Words: attention, multi-tasking, interruption management, multiple resources, time-sharing, display integration, visual scanning, alarms, visual search

Fundamentals

What Is Attention?

Attention may be described as one of a fundamental set of limits to human performance (along with, for example, memory or control precision) on the amount of information that can be processed per unit of time. Of use for the current chapter is the consideration of two metaphors of attention, as a filter and as a fuel (Kramer, Wiegmann, & Kirlik, 2007; Wickens & McCarley, 2008). As a filter, it describes the limits and constraints on the sensory systems (particularly the eyes and ears) to accept and process varying events and elements, up to the level of perception, where the meaning of those events is understood. Thus we conventionally describe the filter metaphor as *selective attention*. As a fuel it describes the limits and constraints on all information processing operations—perception, working memory, decision, and action—to operate concurrently, whether in the service of a single task or in

multitasking. That is, attention characterizes a sort of limited mental energy, fuel, or “resource” that facilitates performance of the relevant process. For example, as the worker “tries harder” to understand a difficult instruction, he or she may lose focus on monitoring other changing variables in the work environment. Thus we can apply the fuel metaphor to *divided attention* between tasks and processes.

Importantly, this dichotomy of metaphors can be broken down by the extent to which the two attention operations succeed or fail. We speak for example of the success of the filter as guiding attention (often our eyes) to relevant sources of information or events in the world; we speak of failures of selective attention as both failures to notice those events at all, and distraction as failures to focus attention on important information as attention is diverted to less important things. We speak of “success” of divided attention, when we can multitask effectively, doing two things at once as well as either

alone. In contrast, failure of divided attention, a matter of degree, ranges from a small dual task decrement in one or the other of two tasks to a complete abandonment of one of them and postponement of its initiation until the other is completed (serial task switching).

What Is Attention in Design?

At a fundamental level, we conceptualize design from a human factors standpoint as an engineering process, whereby the balance between two measurable constructs, performance and workload, is optimized. This balance is complicated in two respects. First, “performance” is itself multifaceted, and in particular in many systems we consider both routine performance and performance in unexpected or “off-nominal” conditions (Burian, 2007; Wickens, Hooy, Gore, Sebok, & Koenicke, 2009). The former is typically the goal of design, but effective human response to off-nominal unexpected conditions depends upon design that supports accurate situation awareness of the task (and the environment in which the task is being performed (Burns et al., 2008; Wickens, 2000a). Such design may not necessarily help routine performance and may sometimes even compromise it. The second complication is that workload should not necessarily be minimized for optimal design, but must be preserved within a range in the middle. This chapter addresses the role of attention in characterizing variables of performance, situation awareness, and workload.

Attention Allocation

As we discuss below, attention may be allocated at two different levels. At the highest level, we can speak of attention—the fuel—as allocated to *tasks*, as tasks may be defined by distinct semi-independent goals (Kirwan & Ainsworth, 1992). Thus the vehicle driver has the task of lane keeping, a second task of navigating, and a third task of dealing with in-vehicle technology (e.g., radio listening, cell phone conversation). Tasks are distinct in this sense in that they usually compete for attentional resources. At the lowest level, we can speak of attention—the filter—as allocated to elements within the environment as well as to internal cognition. Thus, in the vehicle example, the single task of navigation (and higher-level attention directed to the goal of successful navigation) may need to be accomplished by dividing or allocating visual attention (the filter) between a map and the search for landmarks and road signs outside; or between reading a navigation

display, recalling the correct option, and placing the fingers on the correct key for menu choice; or between searching for the road signs and rehearsing the route number to be searched for. In our discussion below, we consider both levels of attention.

A Brief History: Single-Channel Theory and Automaticity

There are two concepts, single-channel processing and automaticity, that are fundamental to most findings and theories in attention, and indeed define endpoints on a sort of continuum from attentional failure to attentional success. Both are deeply rooted in the history of the study of attention (James, 1890; Titchner, 1908).

Single-channel theory (Craik, 1947; Welford, 1967; Pashler, 1998; Broadbent, 1958), the more pessimistic view of human attention, underlies the notion that attention can be focused on only one task at a time, as if performing one task so totally occupies the “single channel” of human cognition and information processing that any other task (usually one arriving later or deemed of lesser importance) must wait, unstarted, until the higher-priority task is completed. Its proponents have cited data in which people must perform two tasks of very high demands at once (like reacting in emergency to an unexpected roadway hazard while dialing a cell phone) or perform two tasks that compete for incompatible resources (like reading a paper document and reading a computer screen).

In stark contrast, the more optimistic view, automaticity (James, 1890; Schneider & Shiffrin, 1977; Fitts & Posner, 1963) defines circumstances when a task requires essentially no attention at all; if it has no attention demands, then ample attention (reserve resources) can be allocated to performing other tasks concurrently without decrement. Classic examples here include walking and talking, or driving (lane keeping) and listening to the radio. In both pairs, the first-mentioned task is so “easy” or automated that it requires little attention.

Single-channel behavior and the perfect time sharing invoked by automaticity of course represent two endpoints on a continuum that can be best defined by the *degree of attentional resources necessary to obtain a given level of performance*. Such a relation between resources and performance is described by the *performance-resource function* (PRF; Norman & Bobrow, 1975), three examples of which are shown in Figure 2.1. The graph line at the bottom (A) suggests a task that would invoke single-channel behavior, since full resources must be allocated to obtain




Figure 2.1 Three examples of the performance-resource function.

perfect performance (or indeed any performance at all). The curve at the top (C) represents an automated task. Perfect performance can be obtained with little or no attention. The graph in the middle (B) highlights the continuum between SCT and automaticity; performance improves up to a point as more resources are allocated to it, but it eventually reaches a level where “trying harder” will not improve performance.

Importantly, the transition from $A \rightarrow B \rightarrow C$ can describe both an intrinsic change in the objective difficulty (complexity or demand value) of the task, or in the subjective difficulty of the task as rendered across three levels of skill development (e.g., novice, journeyman, expert). Important also is the observation that tasks A and C may be performed at equivalent levels in single task conditions; however, when a concurrent task is added, task A will suffer, but C will not.

In the following pages, we describe several general design issues relevant to attention (or attention issues that can be addressed by design)—the role of the filter in noticing, information, access, and search; the role of both the filter and fuel in information integration; the role of the fuel in multitasking that is both parallel and serial; the role of the fuel in mental workload prediction and measurement; and the relationship between workload, situation awareness, and automation. Within each section, we address, where relevant, certain validated *computational models* that can serve the engineering design community.

Noticing and Alerting

Selective attention as the filter can be seen to “tune” toward certain physical events in the environment, while filtering out others. Designers can capitalize on this by assuring that such tuning is focused on important events. Thus a critical design implication of attention is rendered by the attention-capturing properties of alarms and alerts that will direct operators’ attention to events (and locations) that a designer (and sometime

automation) has deemed to be important. The fundamental basis of this approach lies in the fact that people are not very good monitors for infrequent or unexpected events if these are not highlighted in some way, a phenomenon recently described as *change blindness* (Carpenter, 2002; Rensink, 2002; Simons & Levine, 1997; St. John & Smallman, 2008; Wickens Hooley et al., 2009a) or *inattention blindness* (Mack & Rock, 1998). The latter is a form of change blindness that occurs when a change is not noticed, even when directly looking at it.

Alert Salience

Research has identified a number of features of warning systems that will capture attention by making events salient (Boot, Kramer, & Becic, 2007, Iti & Koch, 2000)). For example, appearances of new “objects” in the scene will capture attention, and onsets (increases in luminance) will be more effective in attention capture than will offsets (decreases in luminance or contrast, or disappearing objects; Yantis, 1993). Whether appearing or disappearing, the noticing or attention-capturing properties of these transients is much better when the visual contrast of the changes is larger, the signal/noise ratio is higher (less clutter around the change event location), when visual or cognitive load is lower, and when they occur near or close to foveal vision, than when they are in the periphery (McCarley et al., 2009; Wickens et al., 2009; Steelman-Allen, McCarley, & Wickens, 2011; McKee & Nakayama, 1984). This loss in sensitivity with increasing eccentricity is estimated to be approximately 0.8%/degree (McCarley et al., 2009; Wickens, Alexander, et al., 2003). An extreme example of this eccentric presentation is when the to-be-noticed-event is not in the visual field at all when it occurs (e.g., the eye is closed in a blink or the head is turned beyond about 60 degrees away from the changing element. In these instances, referred to as “completed changes” (Rensink, 2002), change is very hard to notice when fixation is restored to the location of the change.

To some extent, the attention-capturing properties of the physical event (measurable for example by luminance contrast differences) are also modified by knowledge-driven or cognitive processes. One such process is *expectancy*. We will better notice events if they are expected (Wickens, Hooley, et al., 2009); for example, if the operator knows that a system is operating near its limits, he or will more likely expect the warning that those limits have been exceeded and therefore notice the alert when it appears, even if it may not be in foveal vision. A second process

is *tuning*, whereby people are able to “tune” their monitoring to certain event features, to enhance noticing when events contain those features (Most & Astur, 2007; Folk, Remington, & Johnson, 1992, Wolfe & Horowitz, 2004). An obvious case is when the tuned feature is location; people can tune their attention by simply directing their gaze toward the location where an alert is likely to be. But they can also tune attention to be receptive to certain features at a given location: For example, in most cockpit situations, attention is tuned to a red event (e.g., a red light onset) because of the high priority given to red as a warning.

The difference between the *attention-capturing* processes defined by physical elements in the environment (e.g., signal-noise ratio) and the *attention-tuning* processes defined by worker expectations illustrates the more general contrast between what are termed “bottom-up” and “top-down” influences on perception. A final, strong effect on attention capture or noticeability is the ongoing non-visual (auditory and cognitive) workload at the time an event occurs (Fougnie & Marois, 2007).

A computational model called N-SEEV (noticing—salience effort expectancy value; Wickens, Hooy, et al., 2009; Steelman-Allen et al., 2011; Wickens, 2012) can be used to predict the likelihood of detecting an event as a combined function of its salience (Itti & Koch, 2000), expectancy, peripheral eccentricity (from foveal vision), and overall workload. However, in the workplace, as opposed to the laboratory, it is often challenging to determine what the eccentricity of a particular event at a given location may be, as the eyes can be scanning many different locations around the workplace. The SEEV model, the second component of N-SEEV model, predicts the course of this workplace scanning as a context in which the event to be noticed (N) occurs. The SEEV model will be described in a later section.

Beyond the visual modality, there are of course differences in attention-capturing properties between modalities. Most critically, vision is hampered in noticing events in that only about $2 \times 2 = 4$ squared degrees of a momentary visual field of around 60×60 degrees is occupied by foveal vision at any time; that is, only around 0.1%, and noticing degrades rapidly outside of this region. In contrast, events in either the auditory or tactile modality are not much constrained by sensor orientation; they are said to be *omni-directional*, and so auditory (and more recently tactile) warnings have been validated as superior alerts. However, within these non-visual modalities, issues of bottom-up capture (signal-to-

noise ratio) and tuning or expectancy play the same role that they do in vision. As an example, auditory warnings may not be effective in noisy or conversation-rich environments, nor tactile alerts in an environment with extensive physical activity (e.g., a soldier crashing through heavy timber).

Nevertheless, a meta-analysis of noticing events within a visual workplace indicates that the auditory and tactile modality are 15% more effective (faster, more accurate) in capturing attention than are visual interrupting events, even when the latter events are adjacent (in the best case) to the location of the ongoing visual tasks (Wickens, Prinnet et al., 2011; Lu, Wickens et al., 2011; Sarter, this handbook).

Alert Reliability

Most alert systems are imperfect in their reliability. They are designed with algorithms to integrate raw physical data to infer an important or “danger” state (e.g., a malfunction, a fire, or a predicted collision), and if this integrated product exceeds a threshold, the alert activates. However, the raw data are often noisy, and in the case of predictive alerts, circumstances in the environment may change after the alert is given to make the forecast event less likely. The longer this span of prediction is, the lower the reliability. As the obvious consequence, as described by Meyer (2001, 2004) and Meyer and Lee (this handbook), alerts can make one of two types of decision errors: deciding there is not a problem when there is (a “miss”) and deciding that there is a problem when there is not (a “false alert”). When considering the consequences of these two types of errors, most designers quite reasonably assume that misses (or delayed alerts) are worse than false alerts, and so they choose to adjust the threshold lower so that the false alarms are more prevalent. In this case, when the FA rate increases, the system often produces the well-known “cry wolf” problem (Breznitz, 1983; Dixon, Wickens, & McCarley, 2007; Wickens, Rice, et al., 2009; Xiao et al., 2004), whereby operators may turn their attention away from the alerts when they occur and hence are more likely to respond late, or not at all, to true alerts.

Alert Dependence: An Attentional Analysis in a Dual-Task Environment

The effect of alarm reliability can be placed within the broader context of the multitask environment in which alarms are most critical, and the consideration of two cognitive states and two aspects of attention with which those states are associated

(Meyer, 2001, 2004; Meyer & Lee, this handbook; Dixon & Wickens, 2006; Maltz & Shinar, 2003). Thus, in most applications, a busy operator in a multitask environment (driving, flying, health care operations) is *depending* upon the automation to (1) alert him or her if there is a problem, but (2) be “silent” if all is well so that he or she can comfortably turn full attention to the concurrent tasks and away from the domain of the alerted event. As Meyer describes, an operator who responds rapidly to the alerts when they occur is demonstrating the *compliance* to the alert system; one who retains full attention to the concurrent tasks when the alert is silent is demonstrating *reliance* on the alerts. Thus the psychological constructs of compliance and reliance represent two independent aspects of operator dependence upon the alert system (Meyer & Lee, this handbook).

With regard to attention, when the overall reliability of the alert system degrades, both types of automation errors (misses or late alarms and false alarms) may increase. However, a designer-imposed shift in the alert threshold can mitigate the rise in one at the expense of the other. In these cases, data suggest that a rise in false alert rate, with miss rate held constant, will cause a progressive loss in compliance. This “cry wolf” effect can be objectively measured by the response rate, by the response time (to address the alarm), and by a selective attention measure of the time it takes to look at or *switch attention to* the alerting domain (Wickens, Dixon, Goh, & Hammer, 2005). Conversely, an increase in miss rate, with FA rate more or less constant, will lead to a progressive loss in performance on the concurrent task with lower reliance as more attentional resources are reallocated continuously to monitoring the automated domain even when “all is well” (Wickens & Colcombe, 2007). This allocation is directly manifest as increased scanning to any visual display of “raw data” within the alerted domain (Wickens, Dixon, Goh, & Hammer, 2005). These human adjustments in response to failure event frequency may be described as optimal or “eutactic” (Moray & Inagaki, 2000), much as human signal detectors optimally adjust beta in response to signal frequency, as discussed in McCarley & Benjamin (this handbook).

The influences of false alert rate on compliance and miss rate on reliance are not entirely independent in two respects. First, if the threshold of an alert system with constant reliability is varied by the designer, it is obvious that reliance and compliance measures will change in opposite directions. Second, there is some evidence that increasing FA rate not

only degrades compliance but will also degrade reliance (Dixon, Wickens, & McCarley, 2007; Dixon & Wickens, 2006), as if false alarms, being more salient and noticeable than misses, lead to an overall reduction in trust in (and therefore dependence on) the system. So, from the perspective of the impact on human performance in the multitask environment, it appears that FA-prone systems are more problematic than miss-prone (or late-alert-prone) systems. But of course, a full analysis of the appropriate balance between misses and false alarms in alert system design must take into account the primary issue of the costs of overall system misses versus false alerts (i.e., should both the human *and* the alert system miss the dangerous event).

Amplifying and mitigating the AFA problem. Three factors amplify the AFA problem. First, for any given threshold setting, the lower the base rate of events, the greater will be the false alert rate, at least as measured by the proportion of alerts that are false. In some circumstances this can be as high as 0.90. Indeed, in one case (border monitoring for nuclear fuel), it reached 100% (Sanquist, Doctor, & Parasuraman, 2008).

Second, in environments with multiple independent alerts and low thresholds (e.g., the intensive care unit; Seagull et al., 2001), if the probability of a false alert in any given system is even modestly high, then the probability that a single alert within the total workspace will be false can be extraordinarily high. A recent study at a medical center revealed that the typical health care worker was exposed to approximately 850 alerts in a typical workday, many of them undoubtedly false. Nurses experience 841 nuisance alerts/day. Kestin, Miller, and Lockhart (1988) estimated that in the typical operating room an alarm was triggered every 4.5 minutes.

In such circumstances with multiple alarm systems, some of them more prone to false alarms than others, people tend to generalize across the population of all systems, distrusting the good as well as the bad (Keller & Rice, 2010).

Third, the problems with false alarms can obviously be amplified to the extent that those alerts themselves are annoying and intrusive. A visual alert that is false can be fairly effectively “filtered,” since as we noted above, only when it is in the fovea is it most salient. In contrast, the down side of the omni-directionality of auditory or tactile alerts is that the attentional filter cannot restrict access. The increased annoyance accompanying such intrusive false alerts will increase the tendency of workers to deactivate them, or at least try to ignore them (Sorkin, 1989).

Finally, there is emerging evidence that people respond differently when false alerts are clearly “bad” (e.g., the user can obviously perceive that there is no danger) versus when they are “plausible” (e.g., a danger threshold was approached but not quite passed; Lees & Lee, 2007; Wickens, Rice, et al., 2009; see also Madhavan, Wiegmann, & Lacson, 2006). “Cry wolf” behavior is more likely in the former case than in the latter. However, in order for humans to make this determination that a false alarm is plausible, they must be able to monitor the “raw data” independently from, and in parallel with, the automated sensors.

The mitigating solutions for the AFA problem range from the highly intuitive to the less obvious, as we describe below.

- *Increasing alerting system sensitivity in discriminating safe from dangerous conditions.* Often algorithms can be improved and an approach taken over time in developing the airborne traffic alert (TCAS) as designers responded to pilots’ complaints about the high false alarm rate (Rantanen, Wickens, Xu, & Thomas, 2004). An important question in this regard is how low such sensitivity (or reliability) can be before an alerting system becomes no longer effective. One review of alerting studies indicated that with reliabilities above about 0.80 (mean of FA and miss rate), for humans operating in a multitask environment (where attentional resources were at a premium), performance of a human supported by an imperfect alerting system would be better than that of the unaided human (Wickens & Dixon, 2007).

- *Instructing users about the inevitable necessity of some false alarms in uncertain environments, and particularly when the event base rate is lower.* Such instructions can render the false alerts as more “forgivable,” particularly if they are not bad false alerts, as described above.

- *Implementing context sensitive mechanisms* that may raise the threshold during circumstances when the base rate is known to be quite low, and lower it when the base rate is higher (e.g., fire alerts during fire season versus rainy season).

- *Providing the user with rapid (and ideally continuously available) access to the raw data in parallel with the automation.* Hence, to the extent that false alerts are in the “plausible,” not the “bad,” category described above, such access will diminish cry-wolf problems. Indeed, in such a system with raw data access, the activation of the alert may actually reinforce the human’s own raw

data monitoring behavior (if the human detected the pending event before the alert sounded), as well as confirm to the human that the system is in fact well functioning (albeit a little too sensitive). These characteristics appear to have mitigated the “alarm false alarm” issue in some segments of air traffic control (Wickens, Rice, et al., 2009).

- *Developing “likelihood alarms” in which the alert system itself can express its own degree of uncertainty when events occur that are close to the threshold* (Sorkin, Kantowitz, & Kantowitz, 1988; St. Johns & Manes, 2002; Wickens & Colcombe, 2007). Such uncertain-class events can then be associated with a physical sign (e.g., an amber signal) that is less urgent than “sure events” (e.g., red flashing) but more urgent than the sign of “all clear” (e.g., green, or no sign at all). Some evidence suggests that likelihood alerts provide better overall sensitivity than simple two-state alerts (on-off).

- *Informative alerts.* Many complaints about alerts are associated with frustration that, while informing that *something* has gone wrong, they say little about *what* is wrong and *what* to do about it. Such concerns, addressed by making the alerts more informative (e.g., voice alerts), lead us beyond their attention-capturing properties to consideration of the further information properties associated with alerts and other displays, the issue we turn to in the next section.

Attention & Attention Travel in Information Processing Display Layout

Attention, both its filter and fuel capabilities, is particularly challenged in a spatially distributed workspace such as that confronted by the pilot, driver, health care worker, or process controller, where multiple sources of information must be processed as a basis for action and not simply monitored. Such processing may consist of multitasking (as when the driver examines a map while endeavoring to maintain some attention to the roadway), or it may consist of information integration, as when the pilot compares the map with the visual view of landmarks outside the airplane to assure that he or she is on the right track. In such circumstances, we see that attention must *travel* from place to place, an analog to physical travel, and that such travel is not effortless, particularly in a widely distributed visual workspace.

In these circumstances, designers often have an opportunity to “lay out” some aspects of the

workspace to minimize net travel time, according to seven specific principles (Wickens, Vincow, Schopper, & Lincoln, 1997), as we describe in the following. The first two of these principles depend upon defining a “normal line of sight” (NLOS); that is, in a seated workspace, a line about 20 degrees below the horizon extending from the eyes (Sanders & McCormick, 1993). With regard to the point where the line intersects the workspace surface:

1. The most *important* displays should be closer to the NLOS. (This applies particularly to displays whose changes are critical to be noticed in a timely fashion.)
2. The most *frequently used* displays should be closest to the NLOS.
3. Pairs (or N-tuples) of displays used for a single task (i.e., that must be *integrated* or compared and are therefore typically used in sequence) should be close together. In some cases this may involve *database overlay*, as when terrain and weather are superimposed in a pilot’s navigational map so that a safe route through both hazards can be planned (Kroft & Wickens, 2003).
4. Displays *related* to a single class of information should be close together, or *grouped*. This will aid in visual search, as we will see below.
5. Displays should be positioned close to the controls that affect those displays (*display-control compatibility*; Proctor & Proctor, 2006).

We note that in particular, principles 2 (frequency of use) and 3 (relatedness) are designed to minimize the total attention travel time. This optimization, if not followed, may lead to slower performance (since attention travel takes time) but, in a worst case, when attention travel is very effortful, may lead to a relevant display not being visited at all.

Given the role of attention travel in display layout optimization, it is important to realize that travel cost (or *information access cost*) is not a linear function of distance, but instead can be seen to have at least three components (see Wickens, 1993; Wickens & McCarley, 2008): (1) When displays are close together, so that the eye can scan from one to the other without head movement (within about 20 degrees), the cost is minimal and does not change with separation distance. (2) When the displays are separated by more than 20–30 degrees, head movements are required to move the eyes from one to the other, imposing not only a substantially increased cost, but one that grows with the distance (angle) of head movement. (3) Sometimes displays just cannot

be accessed by head movements alone, but rather, require body rotation (checking the blind spot in a car) or, increasingly, key presses or mouse movements to access a particular “page” in a menu or multifunction display. In the latter case, the “distance” of attention travel can be calculated in part by the number of key presses and in part by the cognitive complexity of menu navigation. (e.g., number of options; Seidler & Wickens, 1992, Wickens & Seidler, 1997). Greater information access can not only impose direct time costs but also inhibit information retrieval (Gray & Fu, 2004) and may alter the overall strategy and accuracy of task performance (Morgan, Patrick, et al., 2009).

An important question for designers to answer is what happens when principles “collide” or oppose each other. Suppose, for example, that frequency-of-use dictates that a particular display be close to the NLOS, but integration requires that the same display be close to another, which (for other reasons) has been positioned far from the NLOS. Which principle is more costly to violate? A study that addressed this question had pilots fly with eight different display layouts that either conformed to or violated each of three different principles; frequency of use, integration (sequence of use), and display-control compatibility (Andre & Wickens, 1992). The results revealed that the sequence-of-use principle (close positioning of displays to be integrated for the same task) dominated the frequency-of-use principle, as assessed by overall pilot performance. Both of these dominated display-control compatibility. The impact of these human performance weightings, coupled with others, has been represented in various display layout models summarized in Wickens, Vincow, et al. (1997), which have integrated various elements that influence the efficiency of attention travel, as described above, to provide “figure of merit” estimates of display layout optimization (e.g., Fowler, Williams, Fowler, & Young, 1968).

There are two additional attention-guided principles that can be applied to display layout: A principle of (6) *consistency* dictates that displays should remain in the same consistent location so that they can always be found (selective attention directed there) with minimal interference. Adhering to this principle will not only lead to standardization of layouts across different systems (e.g., aircraft instrument panels always adhere to the basic “T” formation for locating four critical instruments), but adherence will also provide a resistant force against *flexible reconfigurable* display layouts, where designers may choose to reposition displays as a function of work phase (e.g., phase

of flight, or normal vs. abnormal operations), or workers may be given the option of moving displays according to their preference. While such flexibility provides some advantages, these may be offset by the lack of consistency (Andre & Wickens, 1992).

A principle of (7) *clutter avoidance* is one that resists the forces to either put too many displays in a workspace or, in adhering to frequency of use, place all displays tightly clustered or even overlapping. Close proximity achieved via minimizing spatial separation will create clutter—difficulty of focusing attention on individual elements—whenever the spatial separation is less than around 1 degree of visual angle (Broadbent, 1982), and particularly when the elements overlap or are overlaid, as in a HUD display, or a map with text labels overlaying ground features, or overlaying an ATC map (Wickens, 200b).

Head-up displays and head-mounted displays accomplish this by superimposing instruments over an important forward view. The benefit (of not having to move the eyes between the instruments and the forward view) is partially offset by the clutter costs of closely placed information (Wickens, Ververs, & Fadden, 2004). We note here that a special case of close spatial proximity for information to be integrated is represented by *geographical database* overlay; for example, a map of terrain and weather for an aircraft pilot. When the two databases must be integrated (e.g., to find a safe path avoiding both terrain and weather), the close proximity (0 distance) of an overlay provides better performance than a side-by-side presentation of each, despite the greater clutter of the overlay (Kroft & Wickens, 2003; Wickens, 2000b).

The Proximity Compatibility Principle

The theoretical basis for the particular advantage of close proximity displays for information that needs to be integrated (principle 3) lies in the multi-tasking required as the human must retain (often by rehearsal) information from a first-accessed source, while attention travels to the second source for it to be accessed and then compared or combined. At a minimum, the time for travel will degrade memory for the first source. However, if locating the second source requires some search through a cluttered field or (worse yet) accessing another screen via a key press or turning a page, then the mental effort of such access will compete with the retention. This principle, that information that must be integrated in the mind (close mental proximity) should also be close together on a display (close physical proximity),

is referred to as the *proximity compatibility principle* (Wickens & Carswell, 1995; Wickens & McCarley, 2008) and will be addressed further below.

The SEEV Model of Visual Attention Travel

Attention travel across displays and visual workspaces requires eye movements. While in reading text these movements are relatively linear and systematic, in monitoring multi-element displays to supervise dynamic systems, like those of the anesthesiologist, pilot, driver, or process control supervisor, scan paths will be much less predictable. Assisting these predictions is the SEEV model, which was introduced in the previous section in the context of the noticing-SEEV (NSEEV) model of event detection. SEEV predicts steady state scanning around the workspace before the event to be noticed occurs. The integration of its four components—S = salience, E = effort, E = expectancy, and V = value—is based on the prior modeling of Senders (1964, 1980), Sheridan (1970), and Moray (1986), and these are combined additively to predict the distribution of fixation locations. Then, when the to-be-noticed event (TBNE) is scheduled to occur at a specific location in this workspace, SEEV will predict the distribution of eccentricities of that location from the fovea, which in turn predicts the likelihood of detection (diminishing with increasing eccentricity).

The SEEV model has been validated to predict the percentage of time looking at different areas of interest or displays with 80%–90% validity, in workspaces ranging from the live surgical operating table (Koh, Park, Wickens, Teng, & Chia, 2011) to simulations of vehicle driving (Horrey, Wickens, & Consalus, 2006) to both the conventional cockpit (Wickens, Goh, et al., 2003) and the more automated cockpit (Wickens, McCarley, et al., 2008; Steelman-Allen et al., 2011). As noted above, when N is added to SEEV, SEEV then provides the context for predicting eccentricity of the TBNE. N-SEEV has been able to predict pilot detection of a variety of unexpected events both within and outside the cockpit with reasonably high accuracy ($r = 0.75$; Steelman-Allen et al., 2011, Wickens, 2012, Wickens, Hooey, et al., 2009).

The SEEV model predicts how attention is actually allocated across displays. Without the unwanted influence of salience and effort, how attention SHOULD be allocated across displays is defined purely by expectancy (frequency of use and frequency of sequential use) and value. These parameters have been combined in several

computational models of display layout, as discussed above (see Wickens, Vincow, et al., 1997, for review of these).

Display Integration

Design Principles

As noted in the previous section, simply moving displays close together to reduce information access cost can create clutter. There are other means of creating closeness or “proximity” between two or more display elements and hence aid the movement of attention between them, techniques that can loosely be referred to as “display integration.” Many of these are incorporated within the proximity compatibility principle introduced above (see also Wickens & McCarley, 2008). Thus, when spatial proximity cannot be achieved for two elements that are to be integrated (as, for example, when comparing two elements on a map whose coordinates are fixed), the following two techniques can be employed:

- *Linking*, by constructing a physical line between the two, as a line connecting two points on a line graph. Attention can be said to “follow the line,” just as following a road between two geographical locations facilitates the travel from one to another (Jolicoeur & Ingleton, 1991).
- *Common color*, by combining linking and common color. Consider the air traffic control display shown schematically in Figure 2.2 in which planes A and D are at the same altitude and on a collision course. Clearly the controller must mentally integrate the trajectories of the two to determine where and when this collision might take place. Having automation construct a graphic link between them and illuminate them in a distinct color (e.g., red) will facilitate this mental integration computing the anticipated point, time, and separation of closest passage.

The ATC URET conflict alerting system.
Co-altitude (and conflicting) aircraft, need to be mentally integrated to understand conflict. Hence they are **common-colored**




Figure 2.2 Creating proximity in an air traffic display via linking and color.

Besides spatial proximity, linkage, and color, a fourth technique of display integration involves moving two elements so close together that they essentially “fuse” into a single object, a technique known as *object integration*. For example, a single data point on a correlation plot represents two elements, an *X* and a *Y* value (Goettl, Wickens, & Kramer, 1991). The “artificial horizon” on a pilot’s attitude display represents pitch and roll by a single line that can rotate and translate. A single icon object on a weather map may contain several attributes of information. One advantage of object integration, supported by a great deal of research on attention (e.g., Treisman, 1986; Carswell & Wickens, 1996; Duncan, 1984; Scholl, 2001), results because all attributes of a single object are processed more or less in parallel, whereas two separate objects are more likely to be processed in series; hence there is greater efficiency of divided attention between two attributes of the single object display than between two objects.

A fifth technique for display integration, and one that sometimes accompanies object integration, is the creation of *emergent features* (Pomerantz & Pristach, 1989; Bennett & Flach, 2011). This results when multiple elements of a given display “configure” to create a new feature that is not inherent in any of the objects themselves. For example, four bar graphs (e.g., representing engine temperature on four systems) that are all aligned to the same baseline will present an emergent feature of “equality,” which is the co-linearity of their tops, when all are at the same level. Such emergent features can greatly benefit performance to the extent that the feature itself “maps” directly to a critical integration quantity necessary for monitoring and control (Bennett & Flach, 1992; Bennett & Flach, this handbook; Peebles, 2008). If the features are perceptually salient (like the co-linearity above or the symmetrical appearance of certain geometric objects), then direct perception can allow the integration to be achieved without imposing extensive cognitive effort (Vicente, 2002).

Note that the association of object displays with emergent features results because the formation of an object by dimensions, like the length, height, and width of sides and tops of a rectangle, will almost always create emergent features (like the size and shape of the rectangle) that would not exist were the dimensions presented in isolation from each other (e.g., as separate bar graphs; Barnett & Wickens, 1988). However, we also note that if the emergent features of the object are *not* mapped to critical integration task parameters, such object

integration may be of no benefit, and other means of configuring the individual variables may provide better emergent features.

Display Proximity and Clutter

As we have noted above, close proximity achieved via minimizing spatial separation will create clutter. This is one distinct advantage of object integration. Two (or more) attributes of a single object are processed in parallel and hence unlikely to interfere with each other's processing, in contrast to two separate objects occupying the same space (e.g., overlay). Various computational models of clutter have been proposed (e.g., Rosenhotz, Li, & Nakano, 2007; Beck, Lohrenz, & Trafton, 2010).

Extensions of Proximity Compatibility and Object Integration

Two important design concepts related to proximity compatibility are those of *visual momentum* (Woods, 1984; Aretz, 1991; Wickens & McCarley, 2008; Bennett & Flach, 2012) and *ecological interface displays* (Vicente, 2002; Burns & Hajdukiewicz, 2004; Burns et al., 2008). Both have, at their core, the goal of fluently moving attention across complex multi-element workspaces in order to facilitate integration and comparison. Visual momentum is a technique designed to facilitate mental integration of two or more different "views" of a single spatial area or network. For example, one technique of visual momentum would involve presenting a global view of the full workspace, *alongside* a more localized zoom-in view, with the region of the local view highlighted in the global view (Aretz, 1991; Olmos, Liang, & Wickens, 1997; Tang, 2001). Such highlighting allows rapid movement of attention between the two views. A second technique is continuous "panning" rather than abrupt switching between two views of the same region, but from different orientations (Hollands et al., 2008). Visual momentum concepts are particularly valuable when visualizing complex information (Robertson, Czerwinski, et al., 2009; Wickens, Hollands, Banbury, & Parasuraman, 2012).

The concept of an ecological interface is more complex, and space here does not allow much coverage except to note that for very complex systems like power plants, industrial process control industries, or human physiology, there are ways of presenting the multiple variables in such a manner that they directly signal certain critical constraints of the environment or "ecology" that they represent (Burns & Hajdukiewicz, 2004; Burns et al., 2008);

not surprisingly, many of these "ways" capitalize on emergent features and configural displays to graphically represent constraints and boundary conditions in the system (e.g., the balance between mass and energy, or between inflow and outflow, which characterizes stability). Such ecological displays are often found to be most beneficial in fault management, a particular situation when variables must be integrated in new and different ways to diagnose the source of a fault and project its implications to system safety and productivity (Burns, this Handbook, Vicente, 2002; Burns & Hajdukiewicz, 2004).

Visual Search

Visual search is a selective attention function, similar to both noticing and supervisory sampling. However, unlike noticing, search is more goal directed toward finding a predetermined target. In doing so, attention (often coupled with the eyes) usually moves sequentially until the target is found or a decision is made that it is not present (Drury, 2006; Wickens & McCarley, 2008). Search is a key component in many industrial inspection tasks (Drury, 1990, 2006). Thus the primary cognitive demands associated with search *precede* locating the object; whereas the primary task in noticing typically *follows* the triggering event. That said, many variables affect both tasks in the same way: Both usually involve eye movements (when noticing involves a visual event), both are inhibited by a cluttered background and cognitive workload, and both are improved when the target (in search) or the TBNE (in noticing) is salient (flashing, high-contrast, moving, etc.). Importantly, and for a given level of salience, a target will be more likely to be found in a search task than it will be noticed in a noticing task. This difference reflects the added top-down influence of the goal direction of the search task; the search is "tuned" to certain target properties. Both tasks are also influenced by top-down expectancy in other ways. In search, there are two sources of expectancy. Expectancy for target location influences where we look first, and expectancy of whether there is a target at all influences how long we continue a search when the target has not been found (Wolfe, Horowitz, & Kenner, 2005; Drury & Chi, 1995).

From a design perspective, long, tedious searches can have two detrimental influences. First, they can often sacrifice worker efficiency, as, for example, when a computer service worker must spend several seconds searching for a target on a screen, repeating the operation hundreds of times over a workday.

In these circumstances, even milliseconds of added search delay can accumulate large costs (Gray & Boehm-Davis, 2000). Second, they can often inhibit safety, particularly in vehicle control, when long head-down searches (e.g., for a destination on an electronic map) can leave the driver exposed to roadway hazards (Wickens & Horrey, 2009). In another example, analysts computed that the long search time on a railway traffic map spelled the difference between safety and a fatal railway crash, as dispatchers spent 18 precious seconds attempting to locate the source train, causing a flashing collision alert (Stanton & Babar, 2008). This elapsed time spelled the tragic difference between the dispatcher commanding a braking action in time, and too late.

Improving Search

In response concerns such as those described above, a number of attention principles can speak to ways that search can be improved. Some of these solutions include:

- *Target enhancement.* In some circumstances, simple solutions like improving workplace lighting can increase the discriminability between targets and non-targets, a definite advantage when the targets themselves are subtle (like cracks in the hull of an aircraft; Drury, Spencer, & Schurman, 1997).
- *Signal-noise enhancement.* Creative solutions can identify ways to *differentially* amplify the target over the non-targets. For example, if targets are identified by different depths in a three-dimensional display, then providing the user with the ability to change the viewpoint on that display will produce differential motion of targets vs. non-targets (Drury et al., 2001, Drury, 2006).
- *Selective highlighting.* To the extent that the searcher (or another agent) can define features possessed by the target, display technology can then artificially enhance all elements possessing those features—for example, by painting them a different color or increasing their intensity. Thus, for example, in air traffic control, all aircraft flying at a common altitude may be highlighted as particularly relevant because they are more likely to be on a collision course than those at different altitudes (Remington, Johnson, Ruthruff, Gold, & Romera, 2001). Of course, such attention-guidance automation imposes the danger that it could be less than fully reliable (Yeh & Wickens, 2001a; Yeh, Merlo, Wickens, & Brandeberg, 2003; Fisher & Tan, 1989;

Metzger & Parasuraman, 2005). For example, highlighting could be imposed on an element that is not a target, or, more seriously, it could fail to highlight one that is. (These two classes of highlighting errors parallel the two classes of alerting errors discussed previously.) Studies of highlighting validity indicate that people naturally tend to search the highlighted items first (Fisher, Coury, Tengs, & Duffy, 2009), and if there is uncertainty as to whether a target is present or not, people may truncate the search if they fail to find it in the highlighted subset. This behavior would lead to a miss if the target was not highlighted.

- *Search field organization.* In many search fields (e.g., a computer screen), it is possible to impose an organization on the elements to be searched: a linear list or grid. Such organization aids search in two respects. It can help people keep track of examined and not-yet-examined items without excessive burden on memory. It can also avail the opportunity for designers to place the items most likely to be the target of search near the top (for example, the most frequently used items in a computer menu), given the tendency for people to search from top to bottom.

- *Search instructions and target expectancy.* As noted, the expectancy of whether a target is present or not can influence the amount of effort spent on continuing the search when a target is not yet found. Search shows a clear speed-accuracy trade-off, such that longer searches are more likely to turn up a target (Drury, 1994). On the one hand, instructions that emphasize the value of finding the target will produce greater success (but longer search times; Barclay, Vicari, Doughty, Johanson, & Greenlaw, 2006). On the other hand, a low target expectancy will more likely produce a premature termination, leading to a miss (Wolfe et al., 2005). Furthermore, when there may be multiple targets (such as malignant nodules in an x-ray), instructions can counter the tendency to stop the search after a first target is found and impose the search in an exhaustive manner (Barclay et al., 2006).

Modeling Search: The Serial Self-Terminating Model

The serial self-terminating search (STSS) model proposed by Sternberg (1966) is based on data from Neisser (1963) by which attention searches a field of non-targets sometimes containing a target. The

model predicts the time to locate the target or, if it is not present, to decide that it is not. Accordingly, the model predicts that each non-target element is inspected in series, requiring a constant time (T) to decide that each is *not* the target, until the target is reached and a response is made. Thus the search is *self-terminated*. When the target is *not* present, all items must be inspected. When the target *is* present, on average half the items will be inspected. Hence, the slope of the search time as a function of the size of the search field (N) is NT when the target is absent, and $NT/2$ when it is present. Various versions of search models have borrowed from the basic elements of the SSTS model (Drury, 1994; Drury et al., 2001; Teichner & Mocharnuk, 1979; Yeh & Wickens, 2001b; Fisher et al., 1989; Fisher & Tan, 1989; Beck et al., 2010; Nunes, Wickens, & Yin, 2006).

Several modifications and elaborations of this model can be made. For example, if the target is more confusable with the non-targets, T will increase (hence increasing the slope; Geisler & Chou, 1995). If the target is defined by a single salient feature (e.g., red in a sea of green), the slope is essentially 0, describing a *parallel search* process (all items inspected at once). Wolfe (1994; 2007, Wolfe & Horowitz, 2004) has proposed a “guided search” model by which initially several non-targets in the search field can be immediately filtered out (i.e., in parallel), but search through the remainder is serial. This approach has been taken to modeling the benefits of *highlighting* certain key elements of the search field that are assumed to be most relevant, as discussed above (Fisher, Coury, et al., 1989; Beck et al., 2010; Nunes et al., 2006; Yeh & Wickens, 2001b; Wickens, Alexander, et al., 2004).

Attention to Tasks: Multiple Resources

When two tasks must be performed within a narrow window of time, there are two qualitatively different ways in which this can be managed: They can be *time-shared*, wherein the performance of each task is ongoing concurrently, as when listening to a cell phone while driving (Regan, Lee, & Young, 2011; Wickens, Hollands et al., 2012). This is *divided attention* between tasks. Alternatively, they can be performed *in sequence*, as when a driver stops the car before answering the cell phone call. Each situation has very different implications and different sorts of processing operations underlying the success and failure of multitasking, so we consider each in turn.

Concurrent Task Performance: Multiple Resources

According to one prominent theory of multitasking, the multiple resource theory (Navon & Gopher, 1979; Wickens, 1980, 1984, 2002, 2005, 2008a), there are three fundamental elements dictating how well a given task will be performed concurrently with another. First, most intuitively, the difficulty or *attentional resource demand* of both tasks will influence time sharing. Easier tasks (those of lower mental workload, or greater automaticity) will be time shared more effectively (Kahneman, 1973).

Second, a greater degree of *shared versus separate resources* within the human’s information processing structure will increase interference. Wickens (2002) has developed a conception of what those separate resources might be in a way that is consistent with neurophysiological data (Just et al., 2001). For design purposes, these can be broken down in terms of four dichotomies, with “different resources” defined by the two levels of each dichotomy, as follows:

- *processing stages*—perceptual-cognitive (working memory) versus response selection and execution of action
- *processing codes*—spatial versus verbal/linguistic
- *processing modalities (within perception)*—visual versus auditory (and there is now emerging evidence that the tactile channel defines a third perceptual resource category; Lu, Sarter, & Wickens, 2011)
- *visual channels (within visual modality)*—focal (object recognition) versus ambient (motion processing) vision (Previc, 1998, 2000)

Accordingly, as a design and analysis tool (Wickens, 2002, 2005; Wickens, Bagnall, Gosakan, & Walters, 2011), a given task may be defined by levels within one or more of the four dimensions. The interference between two tasks can then be partially predicted by the *number of dimensions on which their demands share common levels*. This prediction of dual task interference is then augmented by summing the total resource demands of the two tasks (independent of their resource competition). A computational version of this model is described in more detail in Wickens, 2005; Sarno and Wickens, 1995; and Wickens, Bagnall et al., 2011.

The third element in predicting success or failure in divided attention between tasks is the *allocation*

policy between them (Norman & Bobrow, 1975; Navon & Gopher, 1979). Intuitively, the more favored task of a pair (the primary task) will preserve its performance close to the single task level, whereas the less favored (the secondary task) will show a greater decrement. This simple feature, allocation policy, describes why the automobile accident rate while using cell phones, while substantial, is not higher than it is: Most drivers still do treat lane keeping and hazard monitoring as a task of higher priority than that of phone conversation.

There is one final factor not accommodated by multiple resource theory that can account for differences in the effectiveness of concurrent task performance, and that is *confusion*, caused by the similarity of elements within the two tasks (Wickens & Hollands, 2000). The more similar those elements are, the more likely there will be cross talk between the two such that, for example, elements of one task show up in the response to the second task. A classic example relates to the challenge of patting your head while rubbing your stomach. Another might be trying to tally or copy student test scores while listening to basketball scores. Note, however, that similarity-based confusion is most likely to occur when the tasks already share some demand for common resources (e.g., in the above two examples, both spatial manual tasks or both auditory/verbal tasks using digits).

Sequential Performance & Task Management

Even when an operator may try to perform two tasks in parallel (albeit with degraded performance on one or both), this may become impossible either because one or both are of high resource demand or because they compete for common incompatible resources, like speaking two different messages at once (the voice can speak only one at a time) or looking to two sources of widely spaced visual inputs. In these circumstances, once the limits of multiple resources have dictated that concurrence is impossible, the first two elements of multiple resource theory (demand and resource structure) no longer play a role in predicting interference. However the third element—allocation policy—now occupies center stage as *the* most important factor in sequential task management: which task is performed and which is completely abandoned or neglected, and for how long.

Two general scenarios underlie the manifestation of sequential task management strategies, both involving a decision process of which task to

perform, and both partially embedded within the framework of queuing theory (Moray, Dessouky, Kijowski, & Adapathya, 1991). One of these is the study of *task switching* (e.g., Rogers & Monsell, 1995; Goodrich, this handbook), and the other is the study of *interruption management* (e.g., Trafton & Monk, 2007). In the former case, the operator is confronted with two tasks and must choose one to initiate first. In the latter case, the operator is already performing one (the “ongoing task”—OT) when a second task (the “interrupting task”—IT) arrives, and must decide whether (or for how long) to continue the OT before switching to the IT, then when to return to the OT. Here researchers often focus on the *quality of OT* upon return (e.g., how fast it is resumed, whether it is resumed where it was “left off,” etc. Trafton & Monk, 2007; Wickens, Hollands et al., 2000).

In both cases, queuing theory can sometimes be applied to determine optimal strategies of task (and interruption) management (Moray et al., 1991; Liao and Moray, 1993). Some of these strategies are quite intuitive, such as when two tasks differ in their importance (or penalty for delayed completion), the more important should be undertaken first. However, when a large number of task features vary between the two, such as their length, their expected duration, their difficulty, the decay of information within a task while it is neglected, or the uncertainty in priority, then assessing optimal solutions becomes very complex. Indeed, in these circumstances it can easily be argued that the mental workload (and time) cost of a human computing the optimal strategy will consume sufficient resources to offset the very goal of trying to make the optimal choice (Raby & Wickens, 1994; Laudeman & Palmer, 1995).

While there are many design-relevant research conclusions in this area, many of these are also based upon only limited data, or data collected in fairly simple laboratory environments. The following paragraphs describe some of the more important of these.

More optimal task switching can be achieved with a *preview* of upcoming tasks (e.g., its duration (Tulga & Sheridan, 1980).

Very slow task switching in multitask environments is suboptimal (Raby & Wickens, 1994), and optimal switching frequency can at least partially be dictated by optimal models (Moray, 1986; Wickens, McCarley, et al., 2008). Particularly in widely distributed visual workspaces, task switching can be partially captured by eye movements, using the

SEEV model to prescribe optimal switching (Kohe et al., 2011).

Very slow task switching characterizes what is sometimes referred to as “attentional tunneling” or “attentional narrowing,” where critical areas of interest (and tasks served by those areas) are neglected for long periods of time, inviting failures to notice key events in those areas (Wickens & Alexander, 2009; Wickens & Horrey, 2009), particularly when those events are unexpected (Wickens, Hooye et al., 2009). In these instances, the “task” that is neglected is often considered the task of maintaining situation awareness (see below).

Three qualitatively different task features tend to induce attentional tunneling, these being extreme levels of interest (such as an engaging cell phone conversation (Horrey, Lesch, & Gabaret, 2009), compelling realistic displays (e.g., a 3-D navigational display; Wickens & Alexander, 2009), and fault management (Moray & Rotenberg, 1989).

Attentional tunneling can be mitigated by salient alarms for neglected tasks (see above), but to be most effective such alarms should be *adaptive* (see Kaber, this handbook), more likely to be activated if automation infers that neglect is taking place (e.g., following an assessment of prolonged head-down orientation in vehicle control).

In interruption management, several variables influence the fluency of task resumption (Dismukes, 2010; Trafton & Monk, 2007; Monk, Trafton, & Boehm-Davis, 2008; Grundgeiger et al., 2010; Smallman & St. John, 2008; Wickens & McCarley, 2008; Morgan, Patrick et al., 2009; Wickens, Hollands et al., 2012), particularly the choice of when to leave an ongoing task (after a subgoal has been completed) and whether a “placeholder” is imposed when the ongoing task is left (e.g., a mark on the page where reading stopped), in order to increase the fluency of return to the OT.

Voice communication tasks tend to be particularly intrusive in interruptions, leading to premature abandonment of ongoing tasks of higher priority (McFarlane & Latorella, 2002; Damos, 1997).

Many aspects of interruption management flow from the study of *prospective memory* (Dismukes, 2010; Loukopoulos, Dismukes, & Barshi, 2009), which is the memory to do a future task. In this particular case, the “future task” is re-engaging the ongoing task following the interruption.

There are beginning to be developed design-oriented solutions that can (a) use automation to monitor the progress of certain types of manual work to assess more appropriate times to interrupt (Bailey &

Konstan, 2006; Dorneich et al., 2012); (b) provide advanced notification of the importance of the interruption so that the operator can decide whether or not to fully abandon the ongoing task or postpone a switch to the interruption task (Ho, Nikolic, Waters, & Sarter, 2004); (c) provide visual placeholders, like a flashing cursor, that will support rapid reacquisition of an ongoing task after the switch (Trafton, Altmann, & Brock, 2005); and (d) provide support tools such as that described by Smallman and St. John, 2008.

Hybrid Models

There is a set of models describing multitasking that are neither strictly parallel (like multiple resources; see above) nor strictly serial (like queuing theory models of sequential performance), but involve scheduling multiple cognitive processes in the service of two tasks that may sometimes be used in series and sometimes in parallel (Meyer & Kieras, 1997; Liu, 1996). One particularly important approach along this line is that of *threaded cognition* (Salvucci & Taatgan, 2008, 2011; Salvucci, this handbook). In particular, the authors have proposed a series of guidelines in the design of multitasking environments.

Conclusion

In conclusion, a great deal of research is required to better understand how people handle sequential tasks under time pressure. One of the more intriguing aspects of this issue involves defining the boundary condition of increasing demands when the multitasker abandons hope of concurrent processing and “regresses” to a sequential mode, ceasing the performance of one task altogether. This “point” is sometimes referred to as a “red line” along a scale of increasing mental workload, imposed by tasks (or sets of tasks), and brings us to the next section on mental workload.

Mental Workload

Mental Workload Assessment

Mental workload may be roughly described as the relation between the attentional resource *demands* (fuel requirements) imposed by tasks and the resources *supplied* by the operator in performing those tasks (fuel available “in the tank”). In the former case, resource requirements can be specified by critical task characteristics that impose greater demands, such as the working memory demands of a task, the number of mental operations, the signal-noise ratio of its displayed elements, the compatibility of mapping from display to control, the precision of required control, the time pressure, or simply the number of tasks imposed at one time. Because a given task

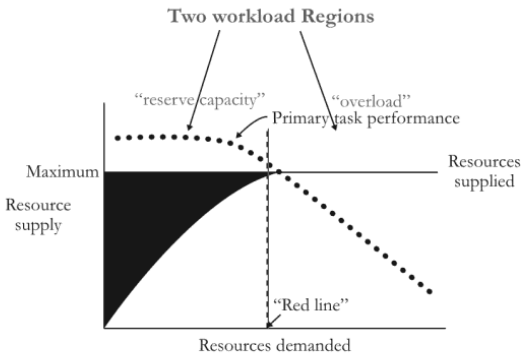


Figure 2.3 The supply-demand curve of resource allocation, illustrating the concept of the “red line”. Wickens, Christopher;

Hollands, Justin G.; *Engineering Psychology and Human Performance*, 3rd Edition, (c) 2000. Reprinted by permission of Pearson Education, Inc., Upper Saddle River, NJ.

environment may be characterized by several of these dimensions at once, each expressed in very different units, the issue of how to combine these into a single metric of “mental workload imposed” is quite challenging, to say the least. It is complicated further because demands of a task configuration will decrease with the skill and practice of the performer.

In the case of resources supplied, there is some evidence that measures of “effort investment” may be more quantifiable, in terms of either physiological measures (Tsang & Vidulich, 2006; Kramer & Parasuraman, 2007) such as heart rate variability or pupil diameter, or in terms of subjective measures (Hart & Staveland, 1988; Hill et al., 1992; Tsang & Vidulich, 2006).

Both measures of resources required and resources supplied (invested) are joined in the “supply-demand” function shown in Figure 2.3, in which increasing demands on the task (x -axis) are met with increasing resources supplied (solid line), up to the point at which resources available are “maxed out.” Performance on the task(s) in question (the dashed line) is perfect up to this point, but further increases in demand cannot be met, and performance then declines. In the parlance of the previous section, this point of inflection on both curves is often referred to as the “red line” of workload, in that designers should strive to maintain task demands always slightly to the left of this point. The desire to stay to the left of the inflection is driven by the design goal of maintaining a margin of “reserve capacity” in order to deal with unexpected emergencies should something go wrong.

In addressing issues of workload, designers are confronted with two top-level questions. First, how can we predict or measure the point along the x -axis of Figure 2.3 imposed by a particular task

requirement in relation to the “red line.” Given the challenges of assessing either resources required or supplied, this can be a difficult enterprise, although progress is being made via the elaborate development of workload assessment measures and computational models of task demand (Laughery, LaBiere, & Archer, 2006). Second, if workload is either predicted or assessed to be above the red line, what can be done to reduce it? Solutions often can be categorized into those that:

- redesign the task (e.g., by changing an interface to use separate resources; by reducing incompatible mappings, by reducing working memory requirements, by facilitating information integration, etc.)
- “redesign” the operator by training
- impose automation

The third solution, using automation to eliminate or reduce human task demands, leads us to a final section relating automation to attention but also invoking a critical third variable, situation awareness.

Attention, Situation Awareness, Workload, and Automation

At a fundamental level, as suggested above, automation and attention demands (workload) are negatively related: The higher the levels of automation that are invoked, the lower the operator workload. The pilot of the modern aircraft with an automated flight management system can fly a complex route with far less hands-on flying than that of a general aviation airplane, where stick, rudder, and throttle may need to be continuously adjusted. But such a simple relationship is complicated in many ways, particularly given the all-important influence of *situation awareness* (Endsley, 1995; Endsley, this handbook; Durso, Rawson, & Giroto, 2007; Banbury & Tremblay, 2004; Parasuraman, Sheridan, & Wickens, 2008; Wickens, 2008b). Thus it is now well established that higher levels of automation will degrade SA in two attention-related respects: monitoring/complacency and working memory.

With regard to monitoring, as automation assumes more tasks that would otherwise require human perception & supervision, the need to monitor what automation is doing decreases. In terms of alerting systems discussed earlier, this was described as increasing *reliance* upon automation (Meyer & Lee, 2004, this handbook), reflected in decreased scanning. Such decreases can be justified as, in some

sense, optimal (Moray, 2003; Moray & Inagaki, 2000), given the low likelihood of automation failure. But if the human supervisor is not looking at automation (or the raw data it is processing), he or she will be slower in noticing those very rare failures in the automated task domain. This is what Endsley has described as a reduction in level 1 situation awareness.

With regard to understanding, the relevant phenomenon in cognitive psychology is referred to as the *generation effect* (Slamecka & Graf, 1978). People are more likely to remember, even briefly, the status of a dynamic system if they have actively responded to change the system than if they have passively witnessed another agent (here automation) making those changes. You remember well the actions you have just taken. The resources invested in making those actions serves you well for future retention. In contrast, decreased memory for (or awareness of) changed state in a highly automated system will leave the monitor of such a system less aware of its precise condition, if a manual takeover is required in a case of a failure. This describes a degradation of Endsley's level 2 SA (understanding); since in many dynamic systems the current state is predictive of future states, it also translates to a degradation of level 3 SA (prediction).

We note then that, as mediated by changes in automation level, there is a direct relationship between SA and workload, a finding that is partially (although imperfectly) documented by empirical research (e.g., Kaber & Endsley, 2004; see Wickens, 2008; Wickens, Li, Santamaria, Sebok, & Sarter, 2010, for a summary). System designers should therefore seek a compromise in adopting a level of automation, between keeping workload manageable and maintaining SA at a sufficiently high level so that the operator can effectively notice and enter the loop should things go wrong.

It is important to realize, however, that the automation-mediated trade-off (between workload and loss of situation awareness) is not inevitable (Tsang & Vidulich, 2006; Wickens, Li, et al., 2010). For example, on the one hand, it may be possible to increase the level of automation to some degree such that workload will decrease but SA will not. This will happen if the curves of SA and WL decrease against automation level increase are non-linear (Wickens, 2008). On the other hand, there are certainly things that can be done to design that will simultaneously reduce workload while improving SA. Certainly training is one: The skilled operator

will have less workload and greater SA than the novice. But importantly, for this chapter, many aspects of *display integration* can also accomplish the combined goals: A well-designed, integrated, and intuitive display can provide a rapid, easy-to-process picture of a dynamic system (supporting situation awareness), and in so doing reduce the cognitive demands of information access, integration, and working memory, simultaneously lowering workload.

Conclusion

In conclusion, we have seen how both the fuel and the filter metaphors provide a useful way of representing many aspects of attention. Derived from basic theory, these two also provide important implications for system design and cognitive engineering. Yet despite the fact that theoretical concepts of attention have been prominent for over a century (James, 1880; Titchner, 1908) and have been applied to system design for over half that time (e.g., Craik, 1947), much remains to be done. For example, the two metaphors need to be better linked to understand the relationship between scanning, selection, and multitasking. In particular, computational models of how attention operates in the complex world beyond the laboratory must be formulated and subjected to rigorous empirical validation, with complex and heterogeneous tasks, to assess the strategies adopted by workers: when to perform tasks concurrently and when, once the red line is exceeded, to abandon and initiate serial multitasking. This is the invitation to the next generation of researchers.

References

- Andre, A. D., & Wickens, C. D. (1992). *Layout analysis for cockpit display systems*. *SID International Symposium Digest of Technical Papers*. Paper presented at Annual Symposium of the Society for Information Display. Seattle Washington
- Aretz, A. J. (1991). The design of electronic map displays. *Human Factors*, 33, 85–101.
- Bailey, B. P., & Konstan, J. A. (2006). On the need for attention-aware systems: Measuring effects of interruption on task performance, error rate, and affective state. *Computers in Human Behavior*, 23, 685–708.
- Banbury, S., & Tremblay, S. (Eds.). (2004). *A cognitive approach to situation awareness: Theory and application*. Aldershot, England: Ashgate.
- Barclay, R. L., Vicari, J. J., Doughty, A. S., Johanson, J. F., & Greenlaw, R. L. (2006). Colonoscopic withdrawal times and adenoma detection during screening colonoscopy. *New England Journal of Medicine*, 355, 2533–2541.
- Barnett, B. J., & Wickens, C. D. (1988). Display proximity in multicue information integration: The benefit of boxes. *Human Factors*, 30, 15–24.

- Beck, R., Lohrenz, M., & Trafton, G. (2010). Measuring search efficiency in complex search tasks. *Journal of Experimental Psychology: Applied*, *16*, 238–250.
- Bennett, K. B., & Flach, J. M. (2011). *Display and interface design: Subtle science, exact art*. Boca Raton, FL: CRC Press.
- Bennett, K. B., & Flach, J. (2012). Visual momentum redux. *International Journal of Human-Computer Studies*, *70*, 399–414.
- Bennett, K. B., & Flach, J. M. (1992). Graphical displays: Implications for divided attention, focused attention, and problem solving. *Human Factors*, *34*, 513–533.
- Boot, W., Kramer, A., & Becic, E. (2007). Capturing attention in the laboratory and the real world. In A. Kramer, D. Wiegmann, & A. Kirlik (Eds.), *Attention: From theory to practice* (pp. 27–45). Oxford, England: Oxford University Press.
- Breznitz, S. (1983). *Cry-wolf: The psychology of false alarms*. Hillsdale, NJ: Erlbaum.
- Broadbent, D. (1958). *Perception and communications*. New York, NY: Pergamon.
- Burian, B. (2007). Perturbing the system: Emergency and off-nominal situations under NextGen. *International Journal of Applied Aviation Studies*, *8*, 114–127.
- Burns, C. M., & Hajdukiewicz, J. R. (2004). *Ecological interface design*. Boca Raton, FL: CRC Press.
- Burns, C. M., Skraaning, G., Jamieson, G. A., Lau, N., Kwok, J., Welch, R., & Andresen, G. (2008). Evaluation of ecological interface design for nuclear process control: Situation awareness effects. *Human Factors*, *50*, 663–679.
- Carpenter, S. (2002). Sights unseen. *APA Monitor*, *32*, 54–57.
- Carswell, C. M., & Wickens, C. D. (1996). Mixing and matching lower-level codes for object displays: Evidence for two sources of proximity compatibility. *Human Factors*, *38*, 1–22.
- Craik, K. W. J. (1947). Theory of the human operator in control systems I: The operator as an engineering system. *British Journal of Psychology*, *38*, 56–61.
- Damos, D. L. (1997). Using interruptions to identify task prioritization in Part 121 air carrier operations. In R. Jensen (Ed.), *Proceedings of the 9th International Symposium on Aviation Psychology*. Columbus, OH: Ohio State University.
- Dismukes, R. K. (2010). Remembrance of things future. In D. Harris (Ed.), *Reviews of human factors & ergonomics* (Vol. 6). Santa Monica, CA: Human Factors & Ergonomics Society.
- Dixon, S. R., & Wickens, C. D. (2006). Automation reliability in unmanned aerial vehicle control: A reliance-compliance model of automation dependence. *Human Factors*, *48*, 474–468.
- Dixon, S. R., Wickens, C. D., & McCarley, J. (2007). On the independence of reliance and compliance: Are automation false alarms worse than misses? *Human Factors*, *49*, 564–573.
- Dorneich, M. C., Ververs, P. M., Mathan, S., Whitlow, S., & Hayes, C. C. (2012). Considering etiquette in the design of an adaptive system. *Journal of Cognitive Engineering and Decision Making*, *6*(2), 243–265.
- Drury, C. G. (1990). Visual search in industrial inspection. In D. Brogan (Ed.), *Visual search* (pp. 263–276). London, England: Taylor & Francis.
- Drury, C. G. (1994). The speed accuracy tradeoff in industry. *Ergonomics*, *37*, 747–763.
- Drury, C. G. (2006). Inspection. In W. Karwowski (Ed.), *International encyclopedia of ergonomics and human factors* (Vol. 2). Boca Raton, FL: Taylor & Francis.
- Drury, C. G., & Chi, C. F. (1995). A test of economic models of stopping policy in visual search. *IIE Transactions*, *27*, 382–393.
- Drury, C., Spencer, F., & Schurman, D. (1997). Measuring human detection performance in aircraft inspection. In *Proceedings of the 41st Annual Meeting of the Human Factors Society*. Santa Monica, CA: Human Factors & Ergonomics Society.
- Drury, C. G., Maheswar, G., Das, A., & Helander, M. G. (2001). Improving visual inspection using binocular rivalry. *International Journal of Production Research*, *39*, 2143–2153.
- Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, *119*, 501–517.
- Durso, F., Rawson, K., & Giroto, S. (2007). Comprehension and situation awareness. In F. Durso (Ed.), *Handbook of applied cognition* (pp. 163–194). Chichester, England: John Wiley.
- Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, *37*, 32–64.
- Fisher, D. L., Coury, B. G., Tengs, T. O., & Duffy, S. A. (1989). Minimizing the time to search visual displays: The role of highlighting. *Human Factors*, *31*(2), 167–182.
- Fisher, J. D., & Tan, K. C. (1989). Visual displays: The highlighting paradox. *Human Factors*, *31*(17–31).
- Fitts, P., & Posner, M. I. *Human performance*. Belmont, CA: Brooks/Cole.
- Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 1030–1044.
- Fougnie, D., & Marois, R. (2007). Executive working memory load induces inattention blindness. *Psychonomic Bulletin & Review*, *14*, 142–147.
- Fowler, R., Williams, W., Fowler, M., & Young, D. (1968). *An investigation of the relationship between operator performance and operator panel layout for continuous tasks* (Technical Report No. 68–170). Wright Patterson AFB, Ohio: US Air Force Flight Dynamics Lab.
- Geisler, W. S., & Chou, K. (1995). Separation of low-level and high-level factors in complex tasks: Visual search. *Psychological Review*, *102*, 356–378.
- Goettl, B. P., Wickens, C. D., & Kramer, A. F. (1991). Integrated displays and the perception of graphical data. *Ergonomics*, *34*, 1047–1063.
- Gray, W., & Boehm-Davis, D. (2000). Milliseconds matter. *Journal of Experimental Psychology: Applied*, *6*, 322–335.
- Gray, W. D., & Fu, W. T. (2004). Soft constraints in interactive behavior: The case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science*, *28*, 359–382.
- Grundgeiger, T., Sanderson, P., Macdougall, H., & Balaubramanian, V. (2010). Interruption management in the intensive care unit. *Journal of Experimental Psychology: Applied*, *16*, 317–334.
- Hart, S. G., & Staveland, L. E. (1988). Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In P. A. H. N. Meshkati (Ed.), *Human mental workload* (pp. 139–183). Amsterdam, The Netherlands: North Holland.
- Hart, S. G., & Wickens, C. D. (2010). Cognitive Workload. *NASA Human Systems Integration handbook*, Chapter

6. Washington, DC: National Aeronautics and Space Administration.
- Hill, S. G., Iavecchia, H., Byers, J., Bittner, A., Zaklad, A., & Christ, R. (1992). Comparison of four subjective workload rating scales. *Human Factors*, *34*, 429–440.
- Ho, C. Y., Nikolic, M. I., Waters, M., & Sarter, N. B. (2004). Not now! Supporting interruption management by indicating the modality and urgency of pending tasks. *Human Factors*, *46*, 399–410.
- Hollands, J. G., Pavlovic, N. J., Enomoto, Y., & Jiang, H. (2008). Smooth rotation of 2-D and 3-D representations of terrain: An investigation into the utility of visual momentum. *Human Factors*, *50*, 62–76.
- Horrey, W. J., Lesch, M. F., & Garabet, A. (2009). Dissociation between driving performance and driver's subjective estimates of performance and workload in dual task conditions. *Journal of Safety Research*, *40*, 7–12.
- Horrey, W. J., Wickens, C. D., & Consalus, K. P. (2006). Modeling drivers' visual attention allocation while interacting with in-vehicle technologies. *Journal of Experimental Psychology: Applied*, *12*(2), 67–86.
- Itti, L., & Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, *40*, 1489–1506.
- James, W. (1890). *Principles of psychology*. New York, NY: Holt. (Reprinted in 1950 by Dover Press.)
- Jolicoeur, P., & Ingleton, M. (1991). Size invariance in curve tracing. *Memory & Cognition*, *19*(1), 21–36.
- Just, M. A., Carpenter, P. A., Keller, T. A., Emery, L., Zajac, H., & Thulborn, K. R. (2001). Interdependence of nonoverlapping cortical systems in dual cognitive tasks. *Neuroimage*, *14*, 417–426.
- Kaber, D. B., & Endsley, M. (2004). The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task. *Theoretical Issues in Ergonomics Science*, *5*, 113–153.
- Kahneman, D. (1973). *Attention and effort*. Englewood Cliffs, NJ: Prentice Hall.
- Kestin, I., Miller, B., & Lockhart, C. (1988). Auditory alarms during anesthesia monitoring. *Anesthesiology*, *69*, 106–109.
- Keller, M. D., & Rice, S. (2010). System-wide versus component-specific trust using multiple aids. *The Journal of General Psychology*, *137*, 114–128.
- Kirwan, B., & Ainsworth, L. (1992). *A guide to task analysis*. London, England: Taylor & Francis.
- Koh, R., Park, T., Wickens, C., Teng O., & Chia, N. (2011). Differences in attentional strategies by novice and experienced operating theatre scrub nurses. *Journal of Experimental Psychology: Applied*, *17*, 233–246.
- Kramer, A. F., & Parasuraman, R. (2007). Neuroergonomics—application of neuroscience to human factors. In J. Caccioppo, L. Tassinari, & G. Berntson (Eds.), *Handbook of psychophysiology* (2nd ed.). New York: Cambridge University Press.
- Kramer, A., Wiegmann, D., & Kirlik, A. (Eds.) (2007). *Attention: from theory to practice*. Oxford, England: Oxford University Press.
- Kroft, P. D., & Wickens, C. D. (2003). Displaying multi-domain graphical database information: An evaluation of scanning, clutter, display size, and user interactivity. *Information Design Journal*, *11*(1), 44–52.
- Laudeman, I. V., & Palmer, E. A. (1995). Quantitative measurement of observed workload in the analysis of aircrew performance. *International Journal of Aviation Psychology*, *5*(2), 187–198.
- Laughery, K. R., LeBiere, C., & Archer, S. (2006). Modeling human performance in complex systems. In G. Salvendy (Ed.), *Handbook of human factors & ergonomics* (pp. 967–996). Hoboken, NJ: John Wiley & Sons.
- Lees, M. N., & Lee, J. D. (2007). The influence of distraction and driving context on driver response to imperfect collision warning systems. *Ergonomics*, *50*, 1264–1286.
- Liao, J., & Moray, N. (1993). A simulation study of human performance deterioration and mental workload. *Le Travail Humain*, *56*(4), 321–344.
- Liu, Y. (1996). Queueing network modeling of elementary mental processes. *Psychological Review*, *103*, 116–136.
- Loukopoulos, L., Dismukes, R. K., & Barshi, I. (2009). *The multi-tasking myth/Handling complexity in real world operations*. Burlington, VT: Ashgate.
- Lu, S., Wickens, C. D., Sarter, N., & Sebok, A. (2011). Informing the design of multimodal displays: A meta-analysis of empirical studies comparing auditory and tactile interruptions. In *Proceedings of the 2011 meeting of the Human Factors & Ergonomics Society*. Santa Monica, CA: Human Factors & Ergonomics Society.
- Mack, A., & Rock, I. (1998). *Inattention blindness*. Cambridge, MA: MIT Press.
- Madhavan, P., Weigmann, D., & Lacson, F. (2006). Automation failures on tasks easily performed by operators undermine trust in automated aids. *Human Factors*, *48*, 241–256.
- Maltz, M., & Shinar, D. (2003). New alternative methods in analyzing human behavior in cued target acquisition. *Human Factors*, *45*, 281–295.
- McCarley, J., Wickens, C., Sebok, A., Steelman-Allen, K., Bzostek, J., & Koenecke, C. (2009). *Control of attention: Modeling the effects of stimulus characteristics, task demands, and individual differences*. University of Illinois Human Factors Division: Urbana, IL: NASA NRA: NNX07AV97A.
- McFarlane, D. C., & Latorella, K. A. (2002). The source and importance of human interruption in human-computer interface design. *Human-Computer Interaction*, *17*, 1–61.
- McKee, S. P., & Nakayama, K. (1984). The detection of motion in the peripheral visual field. *Vision Research*, *24*, 25–32.
- Metzger, U., & Parasuraman, R. (2005). Automation in future air traffic management: Effects of decision aid reliability on controller performance and mental workload. *Human Factors*, *47*, 33–49.
- Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part 1: Basic mechanisms. *Psychological Review*, *104*, 3–65.
- Meyer, J. (2001). Effects of warning validity and proximity on responses to warnings. *Human Factors*, *43*(4), 563–572.
- Meyer, J. (2004). Conceptual issues in the study of dynamic hazard warnings. *Human Factors*, *46*(2), 196–204.
- Monk, C., Trafton, G., & Boehm-Davis, D. (2008). The effect of interruption duration and demand on resuming suspended goals. *Journal of Experimental Psychology: Applied*, *13*, 299–315.
- Moray, N. (1986). Monitoring behavior and supervisory control. In L. K. K. R. Boff & J. P. Thomas (Eds.), *Handbook of perception and performance* (Vol. 2, pp. 40–1–40–51). New York, NY: Wiley & Sons.
- Moray, N. (2003). Monitoring, complacency, scepticism and eutectic behaviour. *International Journal of Industrial Ergonomics*, *31*(3), 175–178.