

Cambridge University Press

978-0-521-70837-1 - Cognitive Science: An Introduction to the Science of the Mind

Jose Luis Bermudez

Frontmatter

[More information](#)

COGNITIVE SCIENCE

An Introduction to the Science of the Mind

This exciting textbook introduces students to the dynamic vibrant area of cognitive science – the scientific study of the mind and cognition. Cognitive science draws upon many academic disciplines, including psychology, computer science, philosophy, linguistics, and neuroscience. This is the first textbook to present a unified view of cognitive science as a discipline in its own right, with a distinctive approach to studying the mind. Students are introduced to the cognitive scientist’s “toolkit” – the vast range of techniques and tools that cognitive scientists can use to study the mind. The book presents the main theoretical models that cognitive scientists are currently using, and shows how those models are being applied to unlock the mysteries of the human mind. *Cognitive Science* is replete with examples, illustrations, and applications and draws on cutting-edge research and new developments to explore both the achievements that cognitive scientists have made, and the challenges that lie ahead.

JOSÉ LUIS BERMÚDEZ is Dean of the College of Liberal Arts and Professor of Philosophy at Texas A&M University. Until 2010 he was Professor of Philosophy and Director of the Philosophy-Neuroscience-Psychology program at Washington University in St. Louis. He has been involved in teaching and research in cognitive science for fifteen years, and is very much involved in bringing an interdisciplinary focus to cognitive science through involvement with conference organization and journals. His 100+ publications include the textbook *Philosophy of Psychology: A Contemporary Introduction* (2005) and a companion collection of readings, *Philosophy of Psychology: Contemporary Readings* (2007). He has authored the monographs *The Paradox of Self-Consciousness* (1998), *Thinking without Words* (2003), and *Decision Theory and Rationality* (2009) in addition to editing a number of collections including *The Body and the Self* (1995), *Reason and Nature* (2002), and *Thought, Reference, and Experience* (2005).

Cambridge University Press

978-0-521-70837-1 - Cognitive Science: An Introduction to the Science of the Mind

Jose Luis Bermudez

Frontmatter

[More information](#)



Cambridge University Press

978-0-521-70837-1 - Cognitive Science: An Introduction to the Science of the Mind

Jose Luis Bermudez

Frontmatter

[More information](#)

COGNITIVE SCIENCE

An Introduction to the Science of the Mind

José Luis Bermúdez



CAMBRIDGE
UNIVERSITY PRESS

Cambridge University Press

978-0-521-70837-1 - Cognitive Science: An Introduction to the Science of the Mind

Jose Luis Bermudez

Frontmatter

[More information](#)

CAMBRIDGE UNIVERSITY PRESS

Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore,
São Paulo, Delhi, Dubai, Tokyo, Mexico City

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org

Information on this title: www.cambridge.org/9780521708371

© José Luis Bermúdez 2010

This publication is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without the written
permission of Cambridge University Press.

First published 2010

Printed in the United Kingdom at the University Press, Cambridge

A catalogue record for this publication is available from the British Library

Library of Congress Cataloguing in Publication data

Bermúdez José Luis.

Cognitive science : an introduction to the science of the mind / José Luis Bermúdez.
p. cm.

Includes bibliographical references.

ISBN 978-0-521-88200-2 – ISBN 978-0-521-70837-1 (pbk.)

1. Cognition. 2. Cognitive science. I. Title.

BF311.B458 2010

153–dc22 2010021896

ISBN 978-0-521-88200-2 Hardback

ISBN 978-0-521-70837-1 Paperback

Additional resources for this publication at www.cambridge.org/bermudez

Cambridge University Press has no responsibility for the persistence or
accuracy of URLs for external or third-party internet websites referred to in
this publication, and does not guarantee that any content on such websites is,
or will remain, accurate or appropriate.

CONTENTS

List of boxes vii
List of figures viii
List of tables xvi
Preface xvii
Acknowledgments xxv

PART I Historical landmarks 2

Introduction to Part I 3
1 The prehistory of cognitive science 5
2 The discipline matures: Three milestones 29
3 The turn to the brain 59

PART II The integration challenge 86

Introduction to Part II 87
4 Cognitive science and the integration challenge 89
5 Tackling the integration challenge 117

PART III Information-processing models of the mind 142

Introduction to Part III 143
6 Physical symbol systems and the language of thought 145
7 Applying the symbolic paradigm 177
8 Neural networks and distributed information processing 215
9 Neural network models of cognitive processes 247

PART IV The organization of the mind 284

Introduction to Part IV 285
10 How are cognitive systems organized? 287

vi Contents

11 Strategies for brain mapping 325
12 A case study: Exploring mindreading 363

PART V New horizons 410

Introduction to Part V 411
13 New horizons: Dynamical systems and situated cognition 413
14 Looking ahead: Challenges and applications 457

Glossary 463
Bibliography 473
Index 486

BOXES

2.1	A conversation with ELIZA	32
3.1	What does each lobe do?	65
3.2	Brain vocabulary	66
4.1	The prisoner's dilemma	107
6.1	Defining well-formed formulas in the propositional logic	149
7.1	Calculating entropy	186
7.2	Calculating information gain	187
7.3	Calculating baseline entropy	189
7.4	Calculating the information gain for Outlook?	190
13.1	Basins of attraction in state space	420

Cambridge University Press

978-0-521-70837-1 - Cognitive Science: An Introduction to the Science of the Mind

Jose Luis Bermudez

Frontmatter

[More information](#)

FIGURES

- 1.1 A rat in a Skinner box
Adapted from Spivey 2007. By permission of Oxford University Press, Inc 8
- 1.2 A 14-unit T-Alley maze
Adapted from Elliott (1928) 9
- 1.3 A cross-maze, as used in Tolman, Ritchie, and Kalish (1946) 11
- 1.4 Schematic representation of a Turing machine
Adapted from Cutland (1980) 15
- 1.5 A sample phrase structure tree for the sentence “John has hit the ball” 18
- 1.6 Donald Broadbent’s 1958 model of selective attention
Adapted by courtesy of Derek Smith 22
- 2.1 A question for SHRDLU about its virtual micro-world
Adapted from Winograd (1972) 33
- 2.2 An algorithm for determining whether a given input is a sentence or not
Adapted from Winograd (1972) 35
- 2.3 Algorithms for identifying noun phrases and verb phrases
Adapted from Winograd (1973) 36
- 2.4 Procedure for applying the concept CLEAR TOP
Adapted from Winograd (1972) 37
- 2.5 SHRDLU acting on the initial command to pick up a big red block
Adapted from Winograd (1972: 8) 38
- 2.6 Instruction 3 in the SHRDLU dialog
Adapted from Winograd (1972: fig. 3) 39
- 2.7 Examples of the three-dimensional figures used in Shepard and Metzler’s 1971 studies of mental rotation.
Adapted from Shepard and Metzler (1971) 42
- 2.8 Examples of vertically and horizontally oriented objects that subjects were asked to visualize in Kosslyn’s 1973 scanning study
Adapted from Kosslyn, Thompson, and Ganis (2006) 46

- 2.9 A table illustrating Marr's three different levels for explaining information-processing systems
From Marr (1982) 48
- 2.10 Two examples of Marr's primal sketch, the first computational stage in his analysis of the early visual system
Adapted from Marr (1982) 51
- 2.11 An example of part of the 2.5D sketch
Adapted from Marr (1982) 52
- 2.12 An illustration of Marr's 3D sketch, showing how the individual components are constructed
Adapted from Marr (1982) 53
- 2.13 The place of the implementational level within Marr's overall theory
Adapted from Marr (1982) 54
- 3.1 The large-scale anatomy of the brain, showing the forebrain, the midbrain, and the hindbrain
Adapted by courtesy of The Huntington's Disease Outreach Project for Education, at Stanford University 63
- 3.2 A vertical slice of the human brain, showing the cerebrum
© TISSUEPIX/SCIENCE PHOTO LIBRARY 64
- 3.3 The division of the left cerebral hemisphere into lobes 64
- 3.4 The primary visual pathway 65
- 3.5 Image showing ventral stream and dorsal stream in the human brain visual system 67
- 3.6 Design and results of Ungerleider and Mishkin's cross-lesion disconnection studies
Adapted from Ungerleider and Mishkin (1982) 69
- 3.7 A generic three-layer connectionist network
Adapted from McLeod, Plunkett, and Rolls (1998) 74
- 3.8 Gorman and Sejnowski's mine/rock detector network
Adapted from Gorman and Sejnowski (1988), printed in Churchland, Paul M., *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*, figure 10.2, page 203 © 1990, Massachusetts Institute of Technology, by permission of the MIT Press. 76
- 3.9 Images showing the different areas of activation (as measured by blood flow) during the four different stages in Petersen *et al.*'s lexical access studies
From Posner and Raichle (1994) 81
- 3.10 A flowchart relating different areas of activation in Petersen *et al.*'s study to different levels of lexical processing
From Petersen *et al.* (1988) 82
- 4.1 Connections among the cognitive sciences, as depicted in the Sloan Foundation's 1978 report
Adapted from Gardner (1985) 91

x List of Figures

- 4.2 Some of the principal branches of scientific psychology 95
- 4.3 Levels of organization and levels of explanation in the nervous system
Adapted from Shepherd (1994) 96
- 4.4 The spatial and temporal resolution of different tools and techniques
in neuroscience
Adapted from Churchland and Sejnowski (1992) 98
- 4.5 The integration challenge and the “space” of contemporary
cognitive science
Adapted by courtesy of David Kaplan 100
- 4.6 A version of the Wason selection task 104
- 4.7 Griggs and Cox’s deontic version of the selection task 105
- 4.8 A microelectrode making an extracellular recording
Reproduced by courtesy of Dwight A. Burkhardt, University
of Minnesota 110
- 4.9 Simultaneous microelectrode and fMRI recordings from a cortical site showing the
neural response to a pulse stimulus of 24 seconds
Adapted from Bandettini and Ungerleider (2001) 112
- 5.1 Two illustrations of the neural damage suffered by the amnesic patient HM
Figure 1, *What’s new with the amnesic patient H.M.?* Nature Neuroscience 2002 Feb.,
3(2): 153–60. 123
- 5.2 Baddeley’s model of working memory 125
- 5.3 The initial stages of a functional decomposition of memory 126
- 5.4 A mechanism for detecting oriented zero-crossing segments
Adapted from Marr and Hildreth (1980) 130
- 6.1 Allen Newell and Herbert Simon studying a search-space
Reproduced by courtesy of Carnegie Mellon University Library 150
- 6.2 A typical traveling salesman problem 151
- 6.3 The structure of Fodor’s argument for the language of thought hypothesis 164
- 6.4 Inside and outside the Chinese room
Courtesy of Robert E. Horn, Visiting Scholar, Stanford University 167
- 7.1 A decision tree illustrating a mortgage expert system
From Friedenbergs and Silverman (2006). 180
- 7.2 A graph illustrating the relation between entropy and probability in the context
of drawing a ball from an urn 185
- 7.3 The first node on the decision tree for the tennis problem 191
- 7.4 The complete decision tree generated by the ID3 algorithm 192
- 7.5 A sample completed questionnaire used as input to an ID3-based
expert system for diagnosing diseases in soybean crops
Adapted from Michalski and Chilauski (1980) 193
- 7.6 Classifying different information-processing models of cognition 195
- 7.7 The basic architecture of WHISPER
From Funt (1980) 196

List of Figures

xi

- 7.8 The starting diagram for the chain reaction problem
From Funt (1980) 198
- 7.9 The result of applying WHISPER's rotation algorithm in order
to work out the trajectory of block B
From Funt (1980) 199
- 7.10 The first solution snapshot output by WHISPER
From Funt (1980) 200
- 7.11 The final snapshot representing WHISPER's solution to the chain reaction problem
From Funt (1980) 201
- 7.12 A map of SHAKEY's physical environment
From Nilsson (1984) 203
- 7.13 A labeled photograph of SHAKEY the robot
Reproduced by courtesy of SRI International, Menlo Park, California 204
- 8.1 Schematic illustration of a typical neuron 219
- 8.2 An artificial neuron 220
- 8.3 Four different activation functions
Adapted from McLeod, Plunkett, and Rolls (1998) 221
- 8.4 Illustration of a mapping function 223
- 8.5 A single-layer network representing the Boolean function AND 225
- 8.6 A single-layer network representing the Boolean function NOT 226
- 8.7 The starting configuration for a single-layer network being trained
to function as a NOT-gate through the perceptron convergence rule 229
- 8.8 Graphical representations of the AND and XOR (exclusive-OR)
functions, showing the linear separability of AND 230
- 8.9 A multilayer network representing the XOR (exclusive-OR) function
Adapted from McLeod, Plunkett, and Rolls (1998) 232
- 8.10 The computational operation performed by a unit in a connectionist model
Adapted from McLeod, Plunkett, and Rolls (1998) 235
- 9.1 Pinker and Prince's dual route model of past tense learning in English 257
- 9.2 Rumelhart and McClelland's model of past tense acquisition
Adapted from Rumelhart, David E., James L. McClelland and PDP Research Group,
Parallel Distributed Processing: Explorations in the Microstructures of Cognition:
Volume 1: Foundations, figure 4, page 242, © 1986 Massachusetts Institute of
Technology, by permission of the MIT Press 258
- 9.3 Performance data for Rumelhart and McClelland's model of past tense learning
Adapted from Rumelhart, David E., James L. McClelland and PDP Research Group,
Parallel Distributed Processing: Explorations in the Microstructures of Cognition:
Volume 1: Foundations, figure 1, page 22, © 1986 Massachusetts Institute of
Technology, by permission of the MIT Press 260
- 9.4 The network developed by Plunkett and Marchman to model children's
learning of the past tense
Adapted from Plunkett and Marchman (1993) 261

Cambridge University Press

978-0-521-70837-1 - Cognitive Science: An Introduction to the Science of the Mind

Jose Luis Bermudez

Frontmatter

[More information](#)

xii List of Figures

- 9.5 A comparison of the errors made by a child and by the Plunkett–Marchman neural network model of tense learning
Adapted from McLeod, Plunkett, and Rolls (1998) 262
- 9.6 Schematic representation of the habituation and test conditions
in Baillargeon’s drawbridge experiments
Baillargeon 1987. Copyright © 1987 by the American Psychological Association.
Reproduced with permission. 264
- 9.7 Schematic representation of an experiment used to test infants’
understanding of Spelke’s principle of cohesion
Adapted from Spelke and Van de Walle (1993) 266
- 9.8 Schematic representation of an experiment testing infants’ understanding
of the principle that only surfaces in contact can move together
Adapted from Spelke and Van de Walle (1993) 267
- 9.9 Schematic depiction of events that accord with, or violate, the continuity
or solidity constraints
Adapted from Spelke and Van de Walle (1993) 268
- 9.10 A series of inputs to the network as a barrier moves in front of a ball
and then back to its original location
Adapted from Munakata, Y., McClelland, J. L., Johnson, M. H.,
Siegler, R. S. (1997)
Copyright © 1997 by the American Psychological Association.
Adapted with permission 272
- 9.11 Recurrent network for learning to anticipate the future position of objects
Adapted from Munakata *et al.* (1997) 273
- 9.12 A balance beam 275
- 9.13 The architecture of the McClelland and Jenkins network for the
balance beam problem
Elman, Jeffrey, Elizabeth Bates, Mark H. Johnson, Annette Karmiloff-Smith,
Domenico Parisi, and Kim Plunkett., *Rethinking Innateness: A Connectionist
Perspective on Development*, figure 3.19, © 1996 Massachusetts Institute
of Technology, by permission of The MIT Press. 276
- 10.1 The architecture of a simple reflex agent
Adapted from Russell and Norvig (2009) 290
- 10.2 The architecture of a goal-based agent
Adapted from Russell and Norvig (2009) 291
- 10.3 The architecture of a learning agent
Russel, Stuart; Norvig, Peter, *Artificial Intelligence: A Modern Approach*,
2nd Edition, © 2003, Pg. 53. Reprinted by permission of Pearson Education,
Inc., Upper Saddle River, NJ 292
- 10.4a Franz Joseph Gall (1758–1828)
Courtesy of Smithsonian Institution Libraries, Washington, DC. 294

List of Figures

xiii

- 10.4b A three-dimensional model of Gall's phrenological map developed by the American phrenologist Lorenzo Niles Fowler (1811–96)
Reproduced courtesy of the Science Museum/Science & Society Picture Library 295
- 10.4c Jerry Fodor (1935–) 296
- 10.5 The evolutionary biologist W. D. Hamilton (1936–2000)
© Jeffrey Joy 308
- 10.6 The ACT-R/PM cognitive architecture
Courtesy of Lorin Hochstein, University of Southern California 316
- 11.1 Luria's 1970 diagram of the functional organization of the brain
Adapted from Luria (1970) 328
- 11.2 Map of the anatomy of the brain showing the four lobes and the Brodmann areas
Reproduced courtesy of Applied Neuroscience Inc. 330
- 11.3 A connectivity matrix for the visual system of the macaque monkey
Adapted from Felleman and Van Essen (1991) 332
- 11.4 An anatomical wiring diagram of the visual system of the macaque monkey
Adapted from Felleman and Van Essen (1991) 333
- 11.5 The results of single-neuron recordings of a mirror neuron in area F5 of the macaque inferior frontal cortex
Adapted from Iacoboni and Dapretto (2006) 337
- 11.6 Typical patterns of EEG waves, together with where/when they are typically found
Courtesy of Jaakko Malmivuo and Robert Plonsey, *Bioelectromagnetism: Principles and Applications of Biomagnetic and Bioelectric Fields*, OUP 1995 338
- 11.7a Common experimental design for neurophysiological studies of attention 343
- 11.7b Example of occipital ERPs recorded in a paradigm of this nature 343
- 11.7c Single-unit responses from area V4 in a similar paradigm 343
- 11.7d Single-unit responses from area V1 showing no effect of attention
Adapted from Luck and Ford (1998), with permission from *Neuroimaging for Human Brain Function* (1998) by the National Academy of Sciences, courtesy of the National Academies Press, Washington DC 344
- 11.8 Frontoparietal cortical network during peripheral visual attention
Gazzaniga, Michael, ed., *The New Cognitive Neurosciences*, Second edition, Plates 30 & 31, © 1999 Massachusetts Institute of Technology, by permission of The MIT Press 349
- 11.9 Peripheral attention vs. spatial working memory vs. saccadic eye movement across studies
Gazzaniga, Michael, ed., *The New Cognitive Neurosciences*, Second edition, Plates 30 & 31, © 1999 Massachusetts Institute of Technology, by permission of The MIT Press 351

xiv List of Figures

- 12.1 An example of metarepresentation 367
- 12.2 The general outlines of Leslie's model of pretend play
Adapted from Leslie (1987) 368
- 12.3 Leslie's Decoupler model of pretense
Adapted from Leslie (1987) 370
- 12.4 The task used by Baron-Cohen, Leslie, and Frith to test for children's
understanding of false belief
Adapted from Baron-Cohen, Leslie, and Frith (1985) 374
- 12.5 The connection between pretend play and success on the false belief task
Adapted from the Open University OpenLearn Unit DSE232_1, courtesy
of the Open University 376
- 12.6 Baron-Cohen's model of the mindreading system 379
- 12.7 What goes on in representing belief 388
- 12.8 What goes on in representing perception 388
- 12.9 A schematic version of standard simulationism
Adapted from Nichols *et al.* (1996) 392
- 12.10 Schematic representation of brain regions associated with the
attribution of mental states
Adapted from Saxe, Carey, and Kanwisher (2004) 397
- 12.11 Schematic overview of the frontoparietal mirror neuron system (MNS)
and its main visual input in the human brain
Adapted from Iacoboni and Dapretto (2006) 403
- 13.1 The trajectory through state space of an idealized swinging pendulum
By permission of M. Casco Associates 416
- 13.2 The state space of a swinging pendulum in a three-dimensional phase space
By permission of M. Casco Associates 417
- 13.3 Illustration of the Watt governor, together with a schematic
representation of how it works
Adapted from Bechtel 1998, *Representations and Cognitive Explanations: Assessing
the Dynamicist's Challenge in Cognitive Science*, © Cognitive Science Society, Inc.,
Vol 22, Issue 3, page 302, figure 2 419
- 13.4 An example of a computational model of motor control
Adapted from Shadmehr and Krakauer (2008) 424
- 13.5 The stage IV search task, which typically gives rise to the
A-not-B-error in infants at around 9 months
Adapted from Bremner (1994) 426
- 13.6 An infant sitting for an A trial (left) and standing for a B trial (right)
Adapted from Smith and Thelen (2003) 428
- 13.7 Applying the dynamical field model to the A-not-B error
Figure 2 in Smith and Thelen, *Development as a Dynamic System*,
Elsevier 2003 430

List of Figures

xv

- 13.8** The organizing principles of biorobotics
Reproduced courtesy of Dimitrios Lambrinos, University of Zurich 436
- 13.9** The anatomy of a cricket, showing the different routes that a
sound can take to each ear
Adapted from Clark (2001) 437
- 13.10** A robot fish called WANDA
Reproduced courtesy of Marc Ziegler, University of Zurich 439
- 13.11** WANDA swimming upwards
From Pfeifer, Iida, and Gómez (2006) 440
- 13.12** Yokoi's robot hand
Reproduced courtesy of Gabriel Gómez, Alejandro Hernandez Arieta,
Hiroshi Yokoi, and Peter Eggenberger Hotz, University of Zurich 441
- 13.13** The Yokoi hand grasping two very different objects
From Pfeifer, Iida, and Gómez (2006) 442
- 13.14** Rodney Brooks's robot Allen, his first subsumption architecture robot
From Brooks (1997) 444
- 13.15** The layers of Allen's subsumption architecture
From Brooks (1997) 445
- 13.16** The Nerd Herd, together with the pucks that they can pick up
with their grippers
Reproduced courtesy of Maja J. Matarić, University of
Southern California 450

TABLES

- 7.1 SHAKEY’S five levels 206
- 7.2 How SHAKEY represents its own state 207
- 7.3 SHAKEY’s intermediate-level actions 208
- 9.1 The stages of past tense learning according to verb type 256
- 10.1 Why we cannot use the language of thought hypothesis
to understand central processing: A summary of Fodor’s worries 303
- 10.2 Comparing the symbolic and subsymbolic dimensions
of knowledge representation in the hybrid ACT-R/PM architecture 320
From Lovett and Anderson (2005)
- 11.1 Comparing techniques for studying connectivity in the brain 340
- 12.1 The three groups studied in Baron-Cohen, Leslie, and Frith 1985 373
- 13.1 The five basis behaviors programmed into Matarić’s Nerd Herd robots 451

PREFACE



About this book

There are few things more fascinating to study than the human mind. And few things that are more difficult to understand. Cognitive science is the enterprise of trying to make sense of this most complex and baffling natural phenomenon.

The very things that make cognitive science so fascinating make it very difficult to study and to teach. Many different disciplines study the mind. Neuroscientists study the mind's biological machinery. Psychologists directly study mental processes such as perception and decision-making. Computer scientists explore how those processes can be simulated and modeled in computers. Evolutionary biologists and anthropologists speculate about how the mind evolved. In fact, there are very few academic areas that are not relevant to the study of the mind in some way. The job of cognitive science is to provide a framework for bringing all these different perspectives together.

This enormous range of information out there about the mind can be overwhelming, both for students and for instructors. I have had direct experience of how challenging this can be, as Director of the Philosophy-Neuroscience-Psychology program at Washington University in St. Louis. The challenge is to give students a broad enough base while at the same time bringing home that cognitive science is a field in its own right, separate and distinct from the disciplines on which it draws. I set out to write this book because my colleagues and I have not yet found a book that really succeeds in doing this.

Different textbooks have approached this challenge in different ways. Some have concentrated on being as comprehensive as possible, with a chapter covering key ideas in each of the relevant disciplines – a chapter on psychology, a chapter on neuroscience, and so on. These books are often written by committee – with each chapter written by an expert in the relevant field. These books can be very valuable, but they

xviii Preface

really give an introduction to the cognitive sciences (in the plural), rather than to cognitive science as an interdisciplinary enterprise.

Other textbook writers take a much more selective approach, introducing cognitive science from the perspective of the disciplines that they know best – from the perspective of philosophy, for example, or of computer science. Again, I have learnt much from these books and they can be very helpful. But I often have the feeling that students need something more general.

This book aims for a balance between these two extremes. Cognitive science has its own problems and its own theories. The book is organized around these. They are all ways of working out the fundamental idea at the heart of cognitive science – which is that the mind is an information processor. What makes cognitive science so rich is that this single basic idea can be (and has been) worked out in many different ways. In presenting these different models of the mind as an information processor I have tried to select as wide a range of examples as possible, in order to give students a sense of cognitive science's breadth and range.

Cognitive science has only been with us for forty or so years. But in that time it has changed a lot. At one time cognitive science was associated with the idea that we can understand the mind without worrying about its biological machinery – we can understand the software without understanding the hardware, to use a popular image. But this is now really a minority view. Neuroscience is now an absolutely fundamental part of cognitive science. Unfortunately this has not really been reflected in textbooks on cognitive science. This book presents a more accurate picture of how central neuroscience is to cognitive science.



How the book is organized

This book is organized into five parts.

Part I: Historical overview

Cognitive science has evolved considerably in its short life. Priorities have changed as new methods have emerged – and some fundamental theoretical assumptions have changed with them. The three chapters in Part I introduce students to some of the highlights in the history of cognitive science. Each chapter is organized around key discoveries and/or theoretical advances.

Part II: The integration challenge

The two chapters in Part II bring out what is distinctive about cognitive science. They do this in terms of what I call the integration challenge. This is the challenge of

developing a unified framework that makes explicit the relations between the different disciplines on which cognitive science draws and the different levels of organization that it studies. In Chapter 4 we look at two examples of *local integration*. The first example explores how evolutionary psychology has been used to explain puzzling data from human decision-making, while the second focuses on what exactly it is that is being studied by techniques of neuro-imaging such as functional magnetic resonance imaging (fMRI).

In Chapter 5 I propose that one way of answering the integration challenge is through developing models of mental architecture. A model of mental architecture includes

- 1 an account of how the mind is organized into different cognitive systems, and
- 2 an account of how information is processed in individual cognitive systems.

This approach to mental architecture sets the agenda for the rest of the book.

Part III: Information-processing models of the mind

The four chapters in Part III explore the two dominant models of information processing in contemporary cognitive science. The first model is associated with the physical symbol system hypothesis originally developed by the computer scientists Allen Newell and Herbert Simon. According to the physical symbol system hypothesis, all information processing involves the manipulation of physical structures that function as symbols. The theoretical case for the physical symbol system hypothesis is discussed in Chapter 6, while Chapter 7 gives three very different examples of research within that paradigm – from data mining, artificial vision, and robotics.

The second model of information processing derives from models of artificial neurons in computational neuroscience and connectionist artificial intelligence. Chapter 8 explores the motivation for this approach and introduces some of the key concepts, while Chapter 9 shows how it can be used to model aspects of language learning and object perception.

Part IV: How is the mind organized?

A mental architecture includes a model both of information processing and of how the mind is organized. The three chapters in Part IV look at different ways of tackling this second problem. Chapter 10 examines the idea that some forms of information processing are carried out by dedicated cognitive modules. It looks also at the radical claim, proposed by evolutionary psychologists, that the mind is simply a collection of specialized modules. In Chapter 11 we look at how some recently developed techniques such as functional neuroimaging can be used to study the organization of the mind. Chapter 12 shows how the theoretical and methodological issues come together by working through an issue that has received much attention in contemporary

xx Preface

cognitive science – the issue of whether there is a dedicated cognitive system response for our understanding of other people (the so-called mindreading system).

Part V: New horizons

As emerges very clearly in the first four parts of the book, cognitive science is built around some very basic theoretical assumptions – and in particular around the assumption that the mind is an information-processing system. In Chapter 13 we look at two ways in which cognitive scientists have proposed extending and moving beyond this basic assumption. One of these research programs is associated with the dynamical systems hypothesis in cognitive science. The second is opened up by the situated/embodied cognition movement.



Using this book in courses

This book has been designed to serve as a self-contained text for a single semester (12–15 weeks) introductory course on cognitive science. Students taking this course may have taken introductory courses in psychology and/or philosophy, but no particular prerequisites are assumed. All the necessary background is provided for a course at the freshman or sophomore level (first or second year). The book could also be used for a more advanced introductory course at the junior or senior level (third or fourth year). In this case the instructor would most likely want to supplement the book with additional readings. There are suggestions on the instructor website (see below).



Text features

I have tried to make this book as user-friendly as possible. Key text features include:

- **Part-openers and chapter overviews**

The book is divided into five parts, as described above. Each part begins with a short introduction to give the reader a broad picture of what lies ahead. Each chapter begins with an overview to orient the reader.

PART I



HISTORICAL LANDMARKS

INTRODUCTION

Here is a short, but accurate, definition of cognitive science: Cognitive science is the science of the mind. Much of this book is devoted to explaining what this means. As with any area of science, cognitive scientists have a set of problems that they are trying to solve and a set of phenomena that they are trying to model and explain. These problems and phenomena are part of what makes cognitive science a distinctive discipline. Equally important, cognitive scientists share a number of basic assumptions about how to go about tackling those problems. They share a very general conception of what the mind is and how it works. The most fundamental driving assumption of cognitive science is that minds are information processors. As we will see, this basic idea can be developed in many different ways, since there are many different ways of thinking about what information is and how it might be processed by the mind.

The chapters in this first section of the book introduce the picture of the mind as an information processor by sketching out some of the key moments in the history of cognitive science. Each chapter is organized around a selection of influential books and articles that illustrate some of the important concepts, tools, and models that we will be looking at in more detail later on in the book. We will see how the basic idea that the mind is an information processor emerged and look at some of the very different ways in which it has been developed.

We begin in Chapter 1 by surveying some of the basic ideas and currents of thought that we can, in retrospect, see as leading into what subsequently emerged as cognitive science. These ideas and currents of thought emerged during the 1930s, 1940s, and 1950s in very different and seemingly unrelated areas. The examples we will look at range from experiments on problem-solving in rats to fundamental breakthroughs in mathematical logic, and from studies of the grammatical structure of language to information-processing models of how input from the senses is processed by the mind.

The early flourishing of cognitive science in the 1960s and 1970s was marked by a series of powerful and influential studies of particular aspects of mental functioning. In Chapter 2 we survey three examples, each of which has been taken by many to be a paradigm of cognitive science in action. These include the studies of mental imagery carried out by Roger Shephard and various collaborators; Terry Winograd's computer program SHRDLU; and David Marr's tri-level model of the early visual system.

The latter decades of the twentieth century saw challenges to some of the basic assumptions of the "founding fathers" of cognitive science. This was cognitive science's "turn to the brain." A crucial factor here was the development of new techniques for studying the brain. These include the possibility of studying the responses of individual neurons, as well as of mapping changing patterns of activation in different brain areas. In Chapter 3 we look at two pioneering sets of experiments. The first is Ungerleider and Mishkin's initial development of the hypothesis that there are two different pathways along which visual information travels through the brain. The second is the elegant use of positron emission tomography (PET) technology by Steve Petersen and collaborators to map how information about individual words is processed in the human brain. Another important factor was the emergence of a new type of model for thinking about cognition, variously known as connectionism or parallel distributed processing. This is also introduced in Chapter 3.

- **Exercises**
These have been inserted at various points within each chapter. They are placed in the flow of the text to encourage the reader to take a break from reading and engage with the material. They are typically straightforward, but for a few I have placed suggested solutions on the instructor website (see below).

32

Box 2.1 A conversation with ELIZA (ELIZA's responses are in upper case)

Men are all alike
IN WHAT WAY?
They are always bugging us about something or other
CAN YOU THINK OF A SPECIFIC EXAMPLE?
Well, my boyfriend made me come here
YOUR BOYFRIEND MADE YOU COME HERE
He says I'm depressed much of the time
I'M SORRY TO HEAR THAT YOU ARE DEPRESSED
It's true. I am unhappy
DO YOU THINK COMING HERE WILL HELP YOU NOT TO BE UNHAPPY?
I need some help. That much seems certain
WHAT WOULD IT MEAN TO YOU IF YOU GOT SOME HELP?
Perhaps I could learn to get along with my mother
TELL ME MORE ABOUT YOUR FAMILY
(From Boden 1977: 106–7)

use the conversations in which it engaged to report on or navigate its environment. So-called *chatbot* programs such as ELIZA are interesting for many reasons, but not as serious models of how we understand and use language.


Q Exercise 2.1 Explain in your own words what you think we can learn from programs such as ELIZA. Is it important that a person might be fooled by ELIZA into thinking that we were communicating with another human being?

Terry Winograd's program SHRDLU, initially presented in his 1970 doctoral dissertation at MIT, was one of the first attempts to write a program that was not just trying to simulate conversation, but that was capable of using language to report on its environment, to plan actions, and to reason about the implications of what is being said to it.

One of the distinctive features of SHRDLU is that it is programmed to deal with a very limited *micro-world* (as opposed to being a general-purpose language program, which is what ELIZA and other chatbot programs are, in their very limited ways). The SHRDLU micro-world is very simple. It consists simply of a number of colored blocks, colored pyramids, and a box, all located on a tabletop, as illustrated in Figure 2.1. (The micro-world is a virtual micro-world, it should be emphasized. Everything takes place on a computer screen.)

2.1 Language and micro-worlds

33



Does the shortest thing the tallest pyramid's support supports support anything green?

Figure 2.1 A question for SHRDLU about its virtual micro-world. (Adapted from Winograd 1972)

SHRDLU is capable of various actions in the micro-world, which it can carry out through a (virtual) robot arm. It can pick up the blocks and pyramids, move them around, and put them in the box. Corresponding to the simplicity of the micro-world, SHRDLU's language is relatively simple. It only has the tools to talk about what is going on in the micro-world.

There are three principal reasons why SHRDLU was very important in the development of cognitive science. The first is that it gave a powerful illustration of how abstract rules and principles such as those in the sort of grammar that we might find in theoretical linguistics could be practically implemented. If we assume that a speaker's understanding of language is best understood as a body of knowledge, then SHRDLU provided a model of how that knowledge could be represented by a cognitive system and how it could be integrated with other, more general, forms of knowledge about the environment.

The second reason for highlighting SHRDLU is that it illustrated the general approach of trying to understand and model cognitive systems by breaking them down into distinct components, each of which carries out a specific information-processing task. One of the many interesting things about SHRDLU is that these distinct components are not completely self-contained. The separate processing systems collaborate in solving information-processing problems. There is *cross-talk* between them, because the programs for each processing system allow it to consult other processing systems at particular moments in the computation.

xxii Preface

- Boxes and optional material
- Boxes have been included to provide further information about the theories and research discussed in the text. Some of the more technical material has been placed in boxes that are marked optional. Readers are encouraged to work through these, but the material is not essential to flow of the text.

186 Applying the symbolic paradigm

Box 7.1 Calculating entropy

OPTIONAL

Entropy in the information-theoretic sense is a way of measuring uncertainty. How do we turn this intuitive idea into a mathematical formula?

To keep things simple we will just calculate the entropy of a set of examples relative to a binary attribute. A binary attribute is one that has two possible values. The example in the text of Black? is a binary attribute, for example. We need some notation – as follows

S	the set of examples
N(S)	the number of examples in S
A	the (binary) attribute
N(A ^(S))	the number of examples with attribute A
N(A ^(S))	the number of examples lacking attribute A

So, the proportion of examples in S with attribute A is given by $\frac{N(A^{(S)})}{N(S)}$ and the proportion of examples in S lacking attribute A is given by $\frac{N(A^{(S)})}{N(S)}$. If we abbreviate these by Prop(A^(S)) and Prop(A^(S)) respectively, then we can calculate the entropy of S relative to A with the following equation

Entropy S/A = – Prop(A^(S)) log₂ Prop(A^(S)) – Prop(A^(S)) log₂ Prop(A^(S))

This is not as bad as it looks! We are working in base 2 logarithms because we are dealing with a binary attribute.

Exercise To make sure that you are comfortable with this equation, refer to the example in the text and check:

(a) that the entropy is 1 when the proportion of black balls is 0.5

(b) that the entropy is 0.88 when the proportion of black balls is 0.7

NB Your calculator may not be able to calculate logarithms to the base 2 directly. The log button will most likely be base 10. You may find the following formula helpful: log₂(x) = log(x) ÷ log(2) for any base.

attribute organizes the remaining examples. It does this by calculating how much the entropy would be reduced if the set were classified according to that attribute. This gives a measure of the information gain for each attribute. Then the algorithm assigns the attribute with the highest information gain to the first node on the tree. Box 7.2 gives the formula for calculating information gain.

Once an attribute has been assigned to the first node we have a tree with at least two branches. And so we have more nodes to which attributes need to be assigned. The algorithm repeats the procedure, starting at the leftmost node. The leftmost node

7.2 ID3: An algorithm for machine learning

187

Box 7.2 Calculating information gain

OPTIONAL

We can measure information gain once we have a way of measuring entropy. Assume that we are starting at a node on the tree. It may be the starting node, but need not be. The node has associated with it a particular subset S* of examples. If the node is the starting node then S* will contain all the examples – i.e. we will have S* = S. If the node is further down the tree then it will be some subset of S – i.e. we have S* ⊂ S.

The first step is to calculate the entropy of S* relative to the target attribute A – i.e. Entropy (S*/A). This can be done using the formula in Box 7.1 and gives the algorithm its baseline again.

Now what we want to do is to calculate how much that uncertainty would be reduced if we had information about whether or not the members of S* have a particular attribute – say B.

So, the second step is to calculate the entropy with respect to the target attribute of the subset of S* that has attribute B – what according to the notation we used in Box 7.1 we call B^(S*). This can be done using the formula from Box 7.1 to give a value for Entropy (B^(S*)/A).

The third step is the same as the second, but in this case we calculate the entropy of B^(S*) with respect to the target attribute – i.e. the subset of S* that does not have attribute B. This gives a value for Entropy (B^(S*)/A).

Finally, the algorithm puts these together to work out the information gain in S* due to attribute B. This is given by the following formula:

Gain (S*, B) = Entropy (S*/A) – Prop (B^(S*)) × Entropy (B^(S*)/A) – Prop (B^(S*)) × Entropy (B^(S*)/A)

As in Box 7.1, Prop (A^(S)) stands for the proportion of S* that has attribute A.

represents a subset S* of the set of examples. So the algorithm calculates the baseline entropy of S* relative to the target attribute. This is the starting point from which it can then calculate which of the remaining attributes has the highest information gain. The attribute with the highest information gain is selected and assigned to the node.

This process is repeated until each branch of the tree ends in a value for the target attribute. This will happen if the attributes on a particular branch end up narrowing the set of examples down so that they all have the same value for the target attribute. When every branch is closed in this way the algorithm halts.

ID3 in action

We can illustrate how ID3 works by showing how it can produce a decision tree for solving a relatively simple problem – deciding whether or not the weather is suitable for playing tennis. In order to apply ID3 we need a database. So imagine that, as keen tennis players who seriously consider playing tennis every day, we collect information

- Summaries, checklists, and further reading
- These can be found at the end of each chapter. The summary shows how the chapter relates to the other chapters in the book. The checklist allows students to review the key points of the chapter, and also serves as a reference point for instructors. Suggestions of additional books and articles are provided to guide students’ further reading on the topics covered in the chapter.

358 Strategies for brain mapping

as telling us about effective connectivity when they are really only telling us about functional connectivity. We must be very careful not to draw conclusions about the causal relations between brain areas and how information flows between them from data that only tell us about correlations between BOLD signal levels in those areas.



SUMMARY

This chapter has continued our exploration of the large-scale organization of the mind. Whereas Chapter 10 focused on issues of modularity, this chapter has looked at some of the ways in which cognitive neuroscience can help us to construct a wiring diagram for the mind. We began by highlighting the complex relations between functional structure and anatomical structure in the brain and then looked at some of the techniques for tracing anatomical connections between different brain areas. Completely different tools are required to move from anatomical connectivity to functional connectivity. We looked at various techniques for mapping the brain through measuring electrical activity and blood flow and blood oxygen levels. These techniques all operate at different degrees of temporal and spatial resolution. As we saw in two case studies, each having to do with a different aspect of the complex phenomenon of attention, mapping the functional structure of the brain requires combining and calibrating different techniques. At the end of the chapter we reviewed some of the pitfalls in interpreting neuroimaging data.



CHECKLIST

It is a basic principle of neuroscience that the cerebral cortex is divided into segregated areas with distinct neuronal populations (the principle of segregation)

- These different regions are distinguished in terms of the types of cell they contain and the density of those cells. This can be studied using staining techniques.
- This anatomical classification of neural areas can serve as a basis for classifying cortical regions according to their function.
- Neuroscientists can study anatomical connectivity (i.e. develop an anatomical wiring diagram of the brain) by using techniques such as tract tracing or diffusion tractography.
- Most of the evidence comes from animal studies. Neuroscientists have developed well worked out models of anatomical connectivity in macaque monkeys, rats, and cats.

Checklist

359

Neuroscientists also adopt the principle of integration – that cognitive functioning involves the coordinated activity of networks of different brain areas

- Identifying these networks requires going beyond anatomical activity by studying what goes on in the brain when it is performing particular tasks.
- Some of the techniques for studying the organization of the mind focus on the brain’s electrical activity. These include electrophysiology, EEG, and MEG.
- These techniques all have high temporal resolution – particularly EEG when it is used to measure ERPs. But the spatial resolution is lower (except for electrophysiology using microelectrodes).
- Other techniques measure blood flow (PET) and levels of blood oxygen (fMRI). These techniques have high spatial resolution, but lower temporal resolution.

The locus of selection problem is the problem of determining whether attention operates early in perceptual processing, or upon representations of objects. It provides a good illustration of how neuroscientists can combine different techniques

- The problem has been studied using EEG to measure ERPs. Attentional effects appear relatively early in the ERP wave following the presentation of a visual stimulus.
- These results can be calibrated with PET studies mapping stages in the ERP wave onto processing in particular brain areas. This calibration reveals attentional effects in areas such as V2 and V4, which carry out very basic processing of perceptual features.
- This resolution of the locus of selection problem seems to be confirmed by single-unit recordings in monkeys.

The locus of selection problem focuses on spatially selective (or visuospatial) attention. Neuroimaging techniques can help identify the neural circuits responsible for attention

- Preliminary evidence from brain-damaged patients (e.g. with hemispatial neglect) points to the involvement of frontal and parietal areas in visuospatial attention.
- This has been confirmed by many experiments on covert attention using PET and fMRI.
- PET and fMRI experiments on humans, together with single-neuron experiments on monkeys, have shown that tasks involving visuospatial attention also generate activation in brain networks responsible for planning motor behavior and for spatial working memory.

The discussion of attention shows that neuroimaging is a very powerful tool for studying cognition. It is not a “window on the mind,” however, and neuroimaging data should be interpreted with caution

- Neuroimaging techniques can only measure cognitive activity indirectly. PET measures blood flow and fMRI measures the BOLD signal. There is a controversy in neuroscience about what type of neural activity is correlated with the BOLD signal (see section 4.5) – and no worked out theory about how that neural activity functions to process information.

Cambridge University Press

978-0-521-70837-1 - Cognitive Science: An Introduction to the Science of the Mind

Jose Luis Bermudez

Frontmatter

[More information](#)

Course website

There is a course website accompanying the book. It can be found at www.cambridge.org/bermudez. This website contains:

- links to useful learning resources, videos, and experimental demonstrations
- links to online versions of relevant papers and online discussions for each chapter
- study questions for each chapter that students can use to structure their reading and that instructors can use for class discussion topics

Instructors can access a password-protected section of the website. This contains:

- sample syllabi for courses of different lengths and different level
- PowerPoint slides
- electronic versions of figures from the text
- suggested solutions for the more challenging exercises and problems

The website is a work in progress. Students and instructors are welcome to contact me with suggestions, revisions, and comments. Contact details are on the website.

CAMBRIDGE

Links to learning resources, videos & experimental demonstrations
 Links to relevant papers & online discussions
 Study questions
FOR INSTRUCTORS
 Sample syllabuses
 Lecture slides
 Suggested solutions
 Figures

Cognitive Science

An Introduction to the Science of the Mind

José Luis Bermúdez

This exciting textbook introduces students to the dynamic area of Cognitive Science – the scientific study of the mind and cognition. Organized thematically instead of by discipline, this is the first textbook to present a unified view of Cognitive Science.

Features:

- Lively, full color presentation with overviews, examples, illustrations, summaries and checklists to aid learning and help students to find and retain information for papers and exams.
- Students are introduced to the cognitive scientist's "toolkit" – the vast range of techniques and tools that cognitive scientists can use to study the mind.
- Draws on cutting-edge research and new developments to explore both the achievements that cognitive scientists have made, and the challenges that lie ahead.

"The text is engaging, well-crafted for an undergraduate audience, and is sure to inspire a generation of students."
 Michael J. Sperry, Professor of Cognitive Science, University of California, Merced

"This book is the most carefully written, thorough, up-to-date, and accessible single-author textbook on the theoretical issues of cognitive science I have read."
 John Dinsart, Department of Philosophy, Rutgers University

Order your inspection copy now | Buy the book

© Cambridge University Press 2010. Privacy Policy | Site map | Contact Us

ACKNOWLEDGMENTS

Many friends and colleagues associated with the Philosophy-Neuroscience-Psychology program at Washington University in St. Louis have commented on sections of this book. I would particularly like to thank Maurizio Corbetta, Frederick Eberhardt, David Kaplan, Clare Palmer, Gualtiero Piccinnini, Marc Raichle, Philip Robbins, David Van Essen, and Jeff Zacks. Josef Perner kindly read a draft of Chapter 12.

I have benefited from the comments of many referees while working on this project. Most remain anonymous, but some have revealed their identity. My thanks to Kirsten Andrews, Gary Bradshaw, Rob Goldstone, Paul Humphreys, and Michael Spivey.

Drafts of this textbook have been used four times to teach PNP 200 Introduction to Cognitive Science here at Washington University in St. Louis – twice by me and once each by David Kaplan and Jake Beck. Feedback from students both inside and outside the classroom was extremely useful. I hope that other instructors who use this text have equally motivated and enthusiastic classes. I would like to record my thanks to the teaching assistants who have worked with me on this course: Juan Montaña, Tim Oakberg, Adam Shriver, and Isaac Wiegman. And also to Kimberly Mount, the PNP administrative assistant, whose help with the figures and preparing the manuscript is greatly appreciated.

A number of students from my Spring 2009 PNP 200 class contributed to the glossary. It was a pleasure to work with Olivia Frosch, Katie Lewis, Juan Manfredi, Eric Potter, and Katie Sadow.

Work on this book has been made much easier by the efforts of the Psychology textbook team at Cambridge University Press – Raihanah Begum, Catherine Flack, Hetty Reid, Sarah Wightman, and Rachel Willsher (as well as to Andy Peart, who signed this book up but has since moved on). They have been very patient and very helpful. My thanks also to Anna Oxbury for her editing and to Liz Davey for coordinating the production process.