Artificial Intelligence and Natural Man

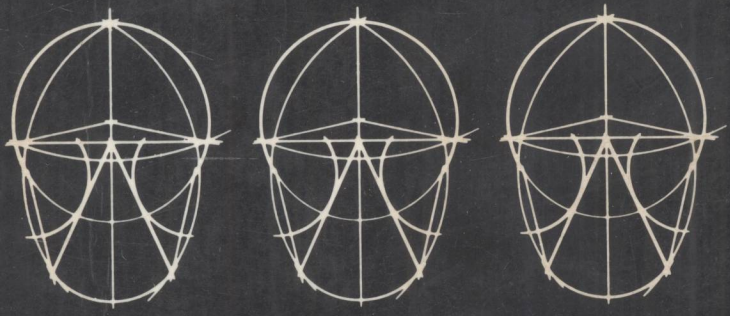

SECOND EDITION, EXPANDED

Margaret A. Boden

BASIC BOOKS, INC.
HARPER TORCHBOOKS TB 5119

TP 1.8
B666
E.2

8863034

ARTIFICIAL INTELLIGENCE AND NATURAL MAN

Second Edition, Expanded

Margaret A. Boden

PROFESSOR OF PHILOSOPHY AND PSYCHOLOGY

UNIVERSITY OF SUSSEX

Basic Books, Inc., Publishers

NEW YORK

Library of Congress Cataloging-in-Publication Data

Boden, Margaret A.
Artificial intelligence and natural man.

Bibliography: p. 501
Includes index.
1. Artificial intelligence. 2. Thought and thinking.
I. Title.
Q335.B56 1987 006.3 86–47739
ISBN 0–465–00456–3 (pbk.)

Copyright © 1977 and 1987 by Margaret A. Boden
Printed in the United States of America
Designed by Vincent Torre
87 88 89 90 HC 9 8 7 6 5 4 3 2 1

ARTIFICIAL INTELLIGENCE AND NATURAL MAN

For Ruskin and Jehane

此为试读，需要完整PDF请访问：www.ertongbook.com

PREFACE TO SECOND EDITION

ARTIFICIAL INTELLIGENCE has grown dramatically since this book first appeared: there are new researchers—and research groups—throughout the industrialized world. Its visibility has increased also, for the media have now discovered it.

But *growth,* whether visible or invisible, can take place without *development.* The recent explosion of funding and publicity is due to commercial and political factors, not to intellectual advances within the field. The central problems of artificial intelligence, and the theoretical basis of its achievements, have remained essentially the same. Naturally, there has been some advance, but most of it is in technological efficiency, not basic scientific understanding.

Despite the passage of time, then, the main message of this book is still timely. The book is not a research text on the nitty-gritty of programming, nor a "list" of the most up-to-date programs around. Indeed, two opening chapters (Chapters 2 and 3) describe a program (a simulation of neurosis) that was already out-of-date when I started writing the first edition in the early 1970s. I chose to discuss it nevertheless—and would do so again today—because its simplicity, and its obvious "stupidity," help introduce two wide-ranging questions: What sort of thing is a computer program? and What difficulties face those who try to make programs as powerful as human thinking? These questions are what this book is about, and they are hardly nearer an answer now than they were a decade ago.

To be sure, if I were to write this book "from scratch" now, it would not be exactly the same. Some recent programs would be included in addition to—occasionally, instead of—the ones originally chosen. And many projects cited in the first edition as exploratory research have been further developed, so could now be described in more detail.

In particular, I would now devote an entire chapter to low-level vision, instead of discussing it within the chapter on psychology (Chapter 13). Low-level vision is the area in which there has been the most significant *theoretical* advance. I would include an account of the current exploratory research on parallel processing ("connectionist") systems. I would allow more space to "production rules," because of their use in commercially available expert systems like those mentioned in Chapter 15. And in view of the current interest in logic programming, I would add PROLOG to the programming languages discussed in Chapter 12. (I discuss the first three

of these topics at length, though from the psychologist's—not the technologist's—point of view, in Boden [in preparation].)

Comparatively minor adjustments would be made also. For example, my discussion of creativity (Chapter 11) would now include programs on heuristic exploration, which have been used to help plan experiments in genetic engineering and design three-dimensional "chips." Chapter 10 would refer to a recent theoretical analysis of learning programs that highlights their computational similarities and differences. Further work on text analysis and story understanding has been done by the workers mentioned in Chapters 7 and 11. And in Chapter 7, too, I would discuss attempts to program *conversation,* which previously I mentioned only in the notes. But none of this new material would invalidate the essential points made in the original version.

For this second edition, then, I have not altered the main text. Instead, I have added a Postscript chapter, and a separate Bibliography listing significant recent work. The Postscript explains the reasons for the recent upsurge of publicity, indicates the areas of theoretical novelty, identifies the main current research issues, and suggests what advances we can—or cannot—expect in the future.

M.A.B.

Brighton, Sussex, April 1986

PREFACE TO FIRST EDITION

ARTIFICIAL INTELLIGENCE is not the study of computers, but of intelligence in thought and action. Computers are its tools, because its theories are expressed as computer programs that enable machines to do things that would require intelligence if done by people.

This book describes artificial intelligence in a way that stresses its human relevance. I have used plain language as far as possible, and have entirely avoided mathematical and formal symbolisms. No specific expertise is presupposed, although readers with psychological or philosophical interests will find many points relating to already familiar issues. In particular, I have not assumed any previous acquaintance with artificial intelligence, nor even any knowledge of programming. So that the reader may gradually develop a sense of what a program is and what a computer can do with it, I have tried to describe programs that do interesting things in a way that allows an understanding of these matters to deepen progressively throughout the book. Later chapters build on earlier ones in the sense that they may continue discussion of examples introduced previously, so the book is best read as a whole, from beginning to end.

I have selected for discussion a number of computer programs likely to be of interest to readers with psychological or philosophical concerns but no programming experience. These programs throw light on the nature of human personality, belief, language and communication, perception, learning, creativity, and problem solving. So that they may provide a starting point from which to progress into detailed study of the programming literature, I have given guidance to further primary sources in the notes. And so that work in artificial intelligence may be related to its wider human context, I have indicated some of the relevant psychological and philosophical literature.

Above all, I have tried to convey a sense of the relevance of artificial intelligence to the understanding of natural man. Contrary to what most people assume, this field of research has a potential for counteracting the dehumanizing influence of natural science, for suggesting solutions to many traditional problems in the philosophy of mind, and for illuminating the hidden complexities of human thinking and personal psychology. The common view that machine research must tend to display us humiliatingly to ourselves as "mere clockwork" is false. The more widely this is realized, the less of a threat will artificial intelligence present to humane conceptions of society.

M.A.B.

Brighton, Sussex
August 1976

ACKNOWLEDGMENTS

I AM deeply grateful to Aaron Sloman for his careful reading of the draft manuscript, and for many conversations on related topics. Helpful comments were provided also by Roy Edgley and Marc Eisenstadt. Much of the book was written during a year's sabbatical leave, granted by the University of Sussex. I am grateful also to my children, to whom this book is dedicated, and who put up with me while I was writing it.

I would also like to thank the following publishers for permission to reprint excerpts from the titles listed below:

Terry Winograd, *Understanding Natural Language* (Academic Press; Edinburgh University Press), pp. 8–15.

Poem by Laurence Lerner from *A.R.T.H.U.R.: The Life and Opinions of a Digital Computer* (University of Massachusetts Press; Harvester Press), pp. 9–10.

Richard Wollheim, ed., *Freud: A Collection of Critical Essays* (Anchor; Doubleday), pp. 242–270.

ACKNOWLEDGMENTS FOR THE SECOND EDITION

I AM grateful to Rudi Lutz for his comments on my Postscript chapter, and to Harold Cohen (of the University of California at San Diego, La Jolla, California) for permission to reproduce two of his computer-generated drawings (photo: Becky Cohen).

"Explain all that," said the Mock Turtle.

"No, no! The adventures first," said the Gryphon in an impatient tone: "explanations take such a dreadful time."

Lewis Carroll

CONTENTS

PREFACE TO SECOND EDITION xi

PREFACE TO FIRST EDITION xiii

ACKNOWLEDGMENTS xiv

Part I

INTRODUCTION

1. *What Is Artificial Intelligence?* 3

PROGRAMS TELL COMPUTERS WHAT TO DO
LANGUAGES THE COMPUTER UNDERSTANDS
COMPUTERS MANIPULATE SYMBOLS

Part II

THE PERSONAL DIMENSION

2. *Artificial Neurosis* 21

OUTLINE DESCRIPTION OF A NEUROTIC PROGRAM
DEFENSE MECHANISMS THAT DISTORT BELIEFS
LATER IMPROVEMENTS

3. *Function and Failure in the Neurotic Program* 34

DATA USED BY THE NEUROTIC PROGRAM
HOW THE NEUROTIC PROGRAM WORKS
FAILINGS OF THE NEUROTIC PROGRAM

4. *Personal Politics and Ideology Machines* 64

SIMULATION OF HOT COGNITION
THE IDEOLOGY MACHINE
A STRUCTURAL THEORY OF BELIEF
BETRAYAL AND OTHER THEMES
SCRIPTS AND SCREENPLAYS

此为试读,需要完整PDF请访问: www.ertongbook.com

Part III

LANGUAGE AND UNDERSTANDING

5. *Responding to Language* 95
PROGRAMS THAT PASS AS PERSONS
HOW PARRY MANAGES IT
HOW ELIZA DIFFERS FROM ELIZA

6. *Intelligence in Understanding* 112
CONVERSATION WITH SHRDLU
HETERARCHICAL THINKING
WHAT SHRDLU KNOWS
ON UNDERSTANDING STORIES

7. *Sense and Semantics* 147
KNOWING WHAT'S GOING TO BE SAID
PROGRAMMING CONCEPTUAL DEPENDENCIES
MACHINE TRANSLATION
LANGUAGE WITHOUT LOGIC

Part IV

THE VISUAL WORLD

8. *Adding the Third Dimension* 179
MODELS, CUES, AND INTERPRETATION
HOW SEE SEES
MAKING SENSE OF PICTURES

9. *Glimpses of the Real World* 207
HALLUCINATION IN THE SERVICE OF TRUTH
MUFFLING THE COMBINATORIAL EXPLOSION
OF MEN AND MOVEMENT

Part V

NEW THOUGHTS FROM OLD

10. *Learning* 247
LEARNING BY EXAMPLE

LEARNING BY BEING TOLD
LEARNING BY DOING

11. *Creativity* 298
CREATING WITHIN A CONVENTION
FICTION AND FRAMES
ANALOGICAL THINKING
CHANGING REPRESENTATIONS

12. *Problem Solving* 345
THE WEAKNESS OF BRUTE FORCE
PLANS FOR PROBLEMS
PROBLEMS AND PROGRAMMING LANGUAGES

Part VI

THE RELEVANCE OF ARTIFICIAL INTELLIGENCE

13. *Psychological Implications* 393
HUMANISM AND MECHANISM IN PSYCHOLOGY
THEORIES AND TESTABILITY
MACHINE VISION AND HUMAN SIGHT

14. *Philosophical Issues* 418
ARTIFICIAL INTELLIGENCE: A CONTRADICTION IN TERMS?
MIND, MECHANISM, AND MACHINES
CAN MACHINE INTELLIGENCE BE ACHIEVED?

15. *Social Significance* 445
PROGRAMS IN PROSPECT
SELVES AND SOCIETY
PRECAUTIONARY MEASURES

16. Postscript 474

BIBLIOGRAPHY 501

BIBLIOGRAPHY TO POSTSCRIPT 521

NOTES 529

INDEX 567

Part I

INTRODUCTION

1

What Is Artificial Intelligence?

ANYONE who mentions artificial intelligence in polite conversation can expect two perplexed reactions: "What?" and "So what?" (Try referring to it as "AI," and you will receive some very strange looks.) Many people have not yet heard of it; few have more than the sketchiest idea of what is involved in it; and most are deeply skeptical of its human value. Even professional conversation can lead to questioning and disagreement about the nature of this novel discipline, for there is no single definition of it that would be judged equally felicitous by all its practitioners.

One thing, however, is certain: artificial intelligence is not the study of computers. Computers are metallic machines of intrinsic interest to electronic engineers but not, as such, to many others. So if you are not enamored of tin cans, you need not fear to meet any in this book.

It would be somewhat more accurate to say that artificial intelligence is the study of computer programs. Indeed, many workers approach it from the context of programming science, being drawn primarily by the challenge it offers to programming techniques. And no one seriously concerned with it can avoid detailed reference to programs and programming.

Yet one leading figure in the field has been heard to retort impatiently to an enthusiastic colleague, "I'm not interested in programs!" This dismissive remark throws light not only on the implicit contrast between

"polite" and "professional" conversation, but also on artificial intelligence itself. For many expert programmers, the activity of programming is strictly subordinate to a wider aim, such as "the development of a systematic theory of intellectual processes, wherever they may be found," or "the study of intelligence as computation."[1] Each of these has been offered as a definition, and each stresses intelligence in general rather than human thought in particular. Other conceptions of artificial intelligence link it more firmly to its human source, emphasizing its potential for generating precise formal models in theoretical psychology.[2]

Still other workers have only a secondary interest (if any) in explaining human thought, concentrating primarily on practical or technological problems such as programming an automatic speech recognizer, filing clerk, or medical diagnostician, or building a robot assembly-line worker or explorer of Mars. People working in a technological context sometimes describe artificial intelligence as "an engineering discipline,"[3] a description hardly calculated to appeal to those with primarily psychological interests.

Clearly, then, the concerns that individuals bring to this type of machine research vary greatly, the preferred definitions differing accordingly. Perhaps the least tendentious definition is M. L. Minsky's—"artificial intelligence is the science of making machines do things that would require intelligence if done by men"—which indicates the general nature of the field without specifying particular aims too closely.[4]

My own interests in artificial intelligence are biased toward its potential for counteracting the dehumanizing influence of natural science, for suggesting solutions to many traditional problems in the philosophy of mind, and for illuminating the hidden complexities of human thinking and personal psychology.

The new concept of "machine" provided by artificial intelligence is so much more powerful than familiar concepts of mechanism that the old metaphysical puzzle of how mind and body can possibly be related is largely resolved. Insights drawn from this source clarify the nature of human purpose, freedom, and moral choice, and help one to understand how it is possible for the character of a human life to depend upon the degree of self-knowledge enjoyed by the person concerned. Artificial intelligence, in short, cannot only acknowledge but can even elucidate the essentially subjective mental realities so stressed by humanist psychologists (as opposed to behaviorists or neurophysiologists).

Equally, it highlights the awesome complexity of the mind, for it leads to an appreciation of the enormous psychological subtlety of "simple" everyday achievements, such as chatting and choosing, of whose computational intricacies one is intuitively quite unaware. Not least, it can suggest specific hypotheses about the thought processes involved, which aid in the development of richly structured theories about the

此为试读,需要完整PDF请访问: www.ertongbook.com