

Cambridge University Press

978-0-521-35944-3 - The Foundations of Artificial Intelligence: A Sourcebook

Edited by Derek Partridge and Yorick Wilks

Frontmatter

[More information](#)

---

## **The foundations of artificial intelligence**

Cambridge University Press

978-0-521-35944-3 - The Foundations of Artificial Intelligence: A Sourcebook

Edited by Derek Partridge and Yorick Wilks

Frontmatter

[More information](#)

---

# The foundations of artificial intelligence

## A sourcebook

Edited by

*Derek Partridge*

Department of Computer Science, University of Exeter  
and

*Yorick Wilks*

Computing Research Laboratory, New Mexico  
State University



Cambridge University Press

978-0-521-35944-3 - The Foundations of Artificial Intelligence: A Sourcebook

Edited by Derek Partridge and Yorick Wilks

Frontmatter

[More information](#)

Published by the Press Syndicate of the University of Cambridge  
The Pitt Building, Trumpington Street, Cambridge CB2 1RP  
40 West 20th Street, New York, NY 10011-4211 USA  
10 Stamford Road, Oakleigh, Victoria 3166, Australia

© Cambridge University Press 1990

First published 1990

Reprinted 1993

Printed in Great Britain at the Athenaeum Press Ltd, Newcastle upon Tyne

*British Library cataloguing in publication data*

The foundations of artificial intelligence: a sourcebook.

I. Artificial intelligence

I. Partridge, D. (Derek), 1945-

II. Wilks, Yorick Alexander

006.3

*Library of Congress cataloguing in publication data*

The foundations of artificial intelligence: a sourcebook /

edited by Derek Partridge and Yorick Wilks.

p. cm.

“Selected papers and subsequent discussion from a NSF-supported workshop on the foundations of AI held in Las Cruces, New Mexico, in 1986” – Pref.

Bibliography.

Includes index.

ISBN 0-521-35103-0. – ISBN 0-521-35944-9 (paperback)

I. Artificial intelligence.

I. Partridge, D. (Derek), 1945- .

II. Wilks, Yorick, 1939- .

0335.F68 1989

006.3 – dc20 89-7321 CIP

ISBN 0 521 35103 0 hardback

ISBN 0 521 35944 9 paperback

Cambridge University Press

978-0-521-35944-3 - The Foundations of Artificial Intelligence: A Sourcebook

Edited by Derek Partridge and Yorick Wilks

Frontmatter

[More information](#)

# Contents

<i>List of contributors</i>	<i>page</i> vii
<i>Preface</i>	ix
<i>Acknowledgments</i>	xiii
<b>1 Introduction</b>	1
What is AI anyway? Roger C. Schank	3
What kind of information processing is intelligence? B. Chandrasekaran	14
<b>2 The formal foundations of AI</b>	47
Non-monotonic reasoning versus logic programming: a new perspective Teodor C. Przymusinski	49
Prolegomena to a theory of mechanized formal reasoning Richard W. Weyhrauch	72
<b>3 Levels of theory</b>	95
AI: a personal view David Marr	97
Has AI helped psychology? Margaret A. Boden	108
What's in an AI program? Derek Partridge	112
<b>4 Programs and theories</b>	119
One small head: models and theories Yorick Wilks	121
The nature of AI principles Alan Bundy and Stellan Ohlsson	135
Artificial methodology meets philosophy Thomas W. Simon	155
<b>5 The role of representations</b>	165
Thinking machines: Can there be? Are we? Terry Winograd	167
Evolution, error, and intentionality Daniel C. Dennett	190
<b>6 The role of programs in AI</b>	213
What kind of field is AI? Alan Bundy	215

Cambridge University Press

978-0-521-35944-3 - The Foundations of Artificial Intelligence: A Sourcebook

Edited by Derek Partridge and Yorick Wilks

Frontmatter

[More information](#)

vi	<i>Contents</i>	
	Programs in the search for intelligent machines: the mistaken foundations of AI Eric Dietrich	223
<b>7</b>	<b>Rational reconstruction as an AI methodology</b>	235
	Three Novelties of AI: theories, programs and rational reconstructions J. A. Campbell	237
	AM: a case study in AI methodology G. D. Ritchie and F. K. Hanna	247
<b>8</b>	<b>Is AI special in regard to its methodology?</b>	267
	Is there anything special about AI? Roger M. Needham	269
	What sort of a thing is an AI experiment? Karen Sparck Jones	274
	We need better standards for AI research John McCarthy	282
<b>9</b>	<b>Does connectionism provide a new paradigm for AI?</b>	287
	Why there STILL has to be a language of thought Jerry A. Fodor	289
	Connectionism and the foundations of AI Paul Smolensky	306
	Some comments on Smolensky and Fodor Yorick Wilks	327
	Representation and high-speed computation in neural networks Paul M. Churchland	337
<b>10</b>	<b>The role of correctness in AI</b>	361
	Does AI have a methodology different from software engineering? Derek Partridge and Yorick Wilks	363
	AI, computer science and education Ranan Banerji	373
<b>11</b>	<b>Limitations on current AI technology</b>	381
	The challenge of open systems Carl Hewitt	383
	Towards a reconciliation of phenomenology and AI Stuart E. Dreyfus and Hubert L. Dreyfus	396
	The superarticulacy phenomenon in the context of software manufacture Donald Michie	411
<b>12</b>	<b>Annotated bibliography on the foundations of AI</b> Imre Balogh and Brian M. Slator	441
	<i>Index of names</i>	492

Cambridge University Press

978-0-521-35944-3 - The Foundations of Artificial Intelligence: A Sourcebook

Edited by Derek Partridge and Yorick Wilks

Frontmatter

[More information](#)

## Contributors

Imre Balogh, Computing Research Laboratory, New Mexico State University

Ranan Banerji, Saint Joseph's University, Philadelphia

Margaret A. Boden, School of Social Sciences, University of Sussex

Alan Bundy, Department of Artificial Intelligence, University of Edinburgh

J. A. Campbell, Department of Computer Science, University College London

B. Chandrasekaran, Laboratory for AI Research, Department of Computer and Information Science, The Ohio State University

Paul M. Churchland, Department of Philosophy, University of California, San Diego

Daniel C. Dennett, Center for Cognitive Studies, Tufts University

Eric Dietrich, State University of New York, Binghamton

Hubert L. Dreyfus, Department of Philosophy, University of California, Berkeley

Stuart E. Dreyfus, Department of Industrial Engineering and Operations Research, University of California, Berkeley

Jerry A. Fodor, Department of Philosophy, City University of New York

F. K. Hanna, Electronics Laboratories, University of Kent

Carl Hewitt, Artificial Intelligence Laboratory, Cambridge, Massachusetts

John McCarthy, Department of Computer Science, Stanford University

The late David Marr

Donald Michie, Machine Intelligence Research Unit, University of Edinburgh

Roger Needham, Computer Laboratory, University of Cambridge

Stellan Ohlsson, Learning Research and Development Center, University of Pittsburgh

Derek Partridge, Department of Computer Science, University of Exeter

Teodor C. Przymusiński, Department of Mathematical Sciences, University of Texas, El Paso

vii

Cambridge University Press

978-0-521-35944-3 - The Foundations of Artificial Intelligence: A Sourcebook

Edited by Derek Partridge and Yorick Wilks

Frontmatter

[More information](#)

---

viii

*Contributors*

G. D. Ritchie, Department of Computer Science, Heriot-Watt University

Roger C. Schank, Department of Computer Science, Yale University

Thomas W. Simon, Department of Philosophy, University of Florida

Brian Slator, Computer Science Department, North Dakota State University

Paul Smolensky, Department of Computer Science and Institute of Cognitive Science, University of Colorado, Boulder

Karen Sparck Jones, Computer Laboratory, University of Cambridge

Richard W. Weyhrauch, Department of Computer Science, Stanford University

Yorick Wilks, Computing Research Laboratory, New Mexico State University

Terry Winograd, Department of Computer Science, Stanford University

Cambridge University Press

978-0-521-35944-3 - The Foundations of Artificial Intelligence: A Sourcebook

Edited by Derek Partridge and Yorick Wilks

Frontmatter

[More information](#)

## Preface

This book collects together a group of fundamental papers on the foundations of artificial intelligence, with selected papers and subsequent discussion from a workshop on the foundations of AI held in Las Cruces, New Mexico, in 1986. The liveliness of the workshop papers and discussions is intended to complement the older, classic papers.

To give the book a structure that will make it accessible to students of the field, we have added a full annotated bibliography, as well as binding explanatory material between the contributions.

At the Las Cruces workshop one of the first questions confronted was the role played by philosophy in the foundations of AI, since philosophy is a subject that comes running whenever foundational or methodological issues arise. The question is important and inevitable but it must always be remembered that there is still an unreformed core of AI practitioners who believe that such assistance – not only from philosophers but from psychologists and linguists as well – can only detract from serious methodological discussions that should take place only between practitioners in private. (A distinguished AI figure said this during the planning period for the workshop.) We need to ask whether that attitude is normal in the sciences, or would-be sciences. For, if it is, then AI is in no way special in these matters, and we can draw benefit from study of the methodologies of other sciences. If not, then we have to ask again whether AI is yet fit to conduct its own methodological discussion. What are the (well-known) precedents and prognoses? Many early methodologists of science were scientists as well as philosophers (Aristotle, Whewell, Kant, Descartes), but many were not (Mill, Pierce). Most in this century were not (Carnap, Popper, Feyerabend, Kuhn, etc.). The more theoretically oriented scientists in this century were not philosophers in any technical sense (e.g. Einstein, Bohr, Planck, Fermi), though some (like Bohm, Feynman, and von Weizsäcker) had a weakness for philosophical expression of what they had to say. They were clearly methodologists of their field, however, but then physics is a very mature and theoretically developed subject.

Perhaps sciences, and even engineering if AI is that, can and should run



Cambridge University Press

978-0-521-35944-3 - The Foundations of Artificial Intelligence: A Sourcebook

Edited by Derek Partridge and Yorick Wilks

Frontmatter

[More information](#)

x

*Preface*

their own methodologies, but only when in that mature state. Until then, they should take all the help, positive or negative, they can get, though this should not be accepted uncritically. A distinguished AI researcher once tried to stop one of the editors from critically questioning a well-known philosopher with the words “You idiot, he’s on our side!”

“Poets have no special insight into what they do or mean” said T. S. Eliot, and there may be some corresponding and essential inadequacy in doing one’s own methodology, wholly autonomously. Perhaps methodologists are a little like accountants, you have to call them in from outside the company to keep straight, at least at intervals. It is this belief that keeps at least a part of philosophy in business, of course.

The fear of the cynic referred to earlier (the one who assumed that discussions might just be wasted if not between AI practitioners) was that the questions the others would want to discuss would not be relevant to the foundations or practice of AI, however much the protagonists might believe them to be: they would always be problem solving in other areas, e.g. clarifying what an “intention” or “mental state” might be, or whether what AI produces are REALLY theories or experiments, in any general sense. Is this fair or true?

The cynic may be right in that no solutions (if there could be any) to such philosophical questions could conceivably affect the agenda of pressing methodological ones that AI has, as seen from the inside. He is unfair in that even partial solutions to such questions (about mental states, intentionality etc.) are important to how AI can describe, present, and interpret what it does, to others and to itself. It has proved almost impossible for AI to describe the function of its programs without drawing on some part of that anthropomorphic vocabulary, and it is the very vocabulary for which AI (and computer science generally) is competing to control, not only with philosophers and psychologists, but poets, novelists, and plain men in the street. It is not because of any particular success of computer science or AI that ordinary people say things like “I think I’m losing core this morning” to mean that their memory is bad.

A serious issue for AI to consider is what its results might be like if purged of all such interpretations: the inverse of the Turing test, as it were, yielding not programs like humans, but successful and useful programs, achieving distinctively human tasks of a high order, but which are described in wholly non-anthropomorphic terms. If that can be done then AI is, at a bound, free, free from discussion by “outsiders” of many of its pretensions and claims. But suppose that cannot be done, and that the chief fault of McDermott’s wonderful squib (D. McDermott, 1981) is that he never saw that.

It is one of the interests of the current connectionist movement in AI, much discussed in this volume, that it is in part an attempt to purge AI of the anthropomorphic imagery that inevitably comes with much of the “symbolic baggage” that it opposes so fiercely. This is not in general a conscious

Cambridge University Press

978-0-521-35944-3 - The Foundations of Artificial Intelligence: A Sourcebook

Edited by Derek Partridge and Yorick Wilks

Frontmatter

[More information](#)*Preface* xi

aim of connectionists, it simply comes with the history: in this case the long association between connectionism, cybernetics, associationism in psychology back to behaviorism and its antipathy to mental talk.

To return: what is this “pressing agenda of internal concerns” of AI referred to above? It contains, and this volume reflects this, such issues as:

- are AI programs distinct from other parts of computer science?
- can there be experiments in AI, as in science, or is their lack evidence that AI is “really engineering”?
- are programs theories, and if not why not?
- is there a special AI form of advance called “throwaway programs”?
- can we isolate parts of computational modeling problems in any determinate way and work on them one at a time?
- why are some levels of program interpretation so obviously to be preferred to others (machine code)?
- why are so few AI “results” repeatable and does it matter?
- why are there no effective statistically significant results in AI and should there be?
- why is it not possible to tell from many AI papers whether a program exists or not?
- why are programs so poorly described?
- is the activity of “rational reconstruction” of the programs of others a defensible and useful one?
- can an AI researcher succeed without having to claim that his program is really intended to model some aspect of human processing?
- are these really pressing methodological issues or just a bunch of family housekeeping worries?
- how does AI combat the claim from computer scientists that it has not contributed anything not already well-known there under other terms?
- is connectionism, as its proponents claim, more experimental/scientific than normal AI, or is the experimentation it offers more social science than science, i.e. producing repeatable results with only that kind of statistical likelihood?

Almost any AI practitioner considers most of those questions important, whether or not they are classified as foundational, philosophical, methodological, or whatever. Compare the concerns of the modern theoretical physicists named above, and how their work is not considered truly methodological by most philosophers of science, but just the scribbles of hopelessly naïve realists.

Why discussion of these issues, whatever they are, is important is that it helps build a critical and reflective tradition in AI, which is perhaps its most striking lack. Having such a tradition may be a necessary precondition for any mature, autonomous, methodology. AI has virtually no tradition of the careful criticism and comparison of ideas, of sifting them for their antecedents or their coherence in a standard scholarly way, so as to extract general principles. Programs and researchers simply come and go endlessly, with no predecessors or successors: techniques are regularly discovered

Cambridge University Press

978-0-521-35944-3 - The Foundations of Artificial Intelligence: A Sourcebook

Edited by Derek Partridge and Yorick Wilks

Frontmatter

[More information](#)

---

xii      *Preface*

every few years (Bundy has documented some of this). Few people seem to mind or notice: after all, creativity and novelty are highly prized in AI and having no memory is one way of ensuring they are always present. All this may be a disbenefit of (the otherwise enormous benefit of) having AI centered in the U.S.: the willingness to try anything once, and not to be hampered by constricting theories and searches for *Anfangspunkte*, and the desire to be positive and uncritical at all costs (except of course when reviewing the grant applications of colleagues: brutal honesty and critical frankness are so important there, particularly because one never has to put one's name to anything).

The aim of this volume, then, is to contribute in a small way towards the establishment of that critical and reflective tradition within AI, and to it, and its saints and martyrs, the book is dedicated.

We would like to thank Jeanine Sandefur for her patience, persistence and sheer hard work when she transformed the various submissions from the contributing authors into a homogeneous and machine-readable whole.

Cambridge University Press

978-0-521-35944-3 - The Foundations of Artificial Intelligence: A Sourcebook

Edited by Derek Partridge and Yorick Wilks

Frontmatter

[More information](#)

## Acknowledgments

The publisher and the editors gratefully acknowledge permission to reproduce the following copyright material:

Daniel Dennett, “Evolution, error and intentionality.” Reprinted with permission, in slightly revised form, from Daniel Dennett: *The Intentional Stance* (MIT Press/Bradford Books 1987).

Jerry A. Fodor, “Why there still has to be a language of thought.” Reprinted with permission from Jerry A. Fodor: *Psychosemantics* (MIT Press/Bradford Books 1987).

Carl Hewitt, “The challenge of open systems.” Reprinted with permission from BYTE magazine (April 1985), pp. 223–42. Copyright © by McGraw-Hill, Inc., New York 10020. All rights reserved.

Alan Bundy and Stellan Ohlsson, “The nature of AI principles.” Reprinted with permission from DAI Research Paper no. 226. The articles on which this research paper was based appeared in the *AISB Quarterly*, nos. 47, 48, 49, and 50.

David Marr, “Artificial intelligence: a personal view.” Reprinted with kind permission of the author’s widow, Lucia Vaina. First published in *Artificial Intelligence*, vol. 9 (1977), pp. 37–48.

Donald Michie, “The superarticulacy phenomenon in the context of software manufacture.” Reprinted with permission from *Proceedings of the Royal Society, London, A*, vol. 405, pp. 185–212.

Derek Partridge and Yorick Wilks, “Does AI have a methodology which is different from software engineering.” Reprinted with permission from *AI Review*, vol. 1, no. 2 (1987), pp. 111–20. Copyright © Intellect Ltd.

Cambridge University Press

978-0-521-35944-3 - The Foundations of Artificial Intelligence: A Sourcebook

Edited by Derek Partridge and Yorick Wilks

Frontmatter

[More information](#)

---

xiv      *Acknowledgments*

Graeme Ritchie and F. K. Hanna, “AM: a case study in AI methodology.” Reprinted with permission from *AI Journal*, vol. 23, no. 3 (1984), pp. 249–63. Copyright © Elsevier Science Publishers BV.

Paul Smolensky, “On the proper treatment of connectionism.” Reprinted with permission, in shortened form, from *Behavioral and Brain Sciences*, vol. 11, no. 1, March 1988. Copyright © Cambridge University Press.

Richard W. Weyhrauch, “Prolegomena to a theory of mechanized formal reasoning.” Reprinted with permission, in shortened and revised form, from *AI Journal*, vol. 12, no. 1–2 (1980), pp. 133–70. Copyright © Elsevier Science Publishers BV.

Terry Winograd, “Thinking machines: Can there be? Are we?” Reprinted with permission of the Stanford Humanities Center, Stanford University.