Formal Methods in Artificial Intelligence

Cambridge Tracts in Theoretical Computer Science

Managing Editor: Professor C. J. van Rijsbergen, Computing Science Department, University of Glasgow

Editorial Board:
S. Abramsky, Department of Computing Science, Imperial College of Science and Technology, London
P. H. Aczel, Department of Computer Science, University of Manchester
J. W. de Bakker, Centrum voor Wiskunde en Informatica, Amsterdam
J. A. Goguen, Programming Research Group, University of Oxford
J. V. Tucker, Department of Mathematics and Computer Science, University College of Swansea

Titles in the Series

1. G. J. Chaitin *Algorithmic Information Theory*
2. L. C. Paulson *Logic and Computation*
3. J. M. Spivey *Understanding Z*
4. G. E. Revesz *Lambda Calculus, Combinators and Functional Programming*
5. S. Vickers *Topology via Logic*
6. A. M. Ramsay *Formal Methods in Artificial Intelligence*
7. J-Y. Girard, Y. Lafont & P. Taylor *Proofs and Types*
8. J. Clifford *Formal Semantics & Pragmatics for Natural Language Processing*
9. M. Winslett *Updating Logical Databases*
10. K. McEvoy & J. V. Tucker (eds.) *Theoretical Foundations of VLSI Design*
11. T. H. Tse *A Unifying Framework for Structured Analysis and Design Models*
12. G. Brewka *Nonmonotonic Reasoning*
13. G. Smolka *Logic Programming over Polymorphically Order-Sorted Types*
15. S. Dasgupta *Design Theory and Computer Science*
17. J. C. M. Baeten (ed.) *Applications of Process Algebra*
18. J. C. M. Baeten & W. P. Weijland *Process Algebra*

FORMAL METHODS IN ARTIFICIAL INTELLIGENCE

ALLAN RAMSAY
Department of Computer Science, University College, Dublin

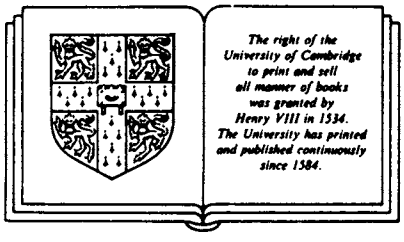

CAMBRIDGE UNIVERSITY PRESS
Cambridge
New York Port Chester
Melbourne Sydney

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo, Delhi

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/9780521424219

First published 1988
Reprinted 1989
First paperback edition 1991
Re-issued in this digitally printed version 2009

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

Library of Congress Cataloguing in Publication data
Ramsay, Allan, 1953 –
Formal Methods in Artificial Intelligence.
(Cambridge Tracts in Theoretical Computer Science; 6)
Bilbiography: p.
Includes index.
1. Artificial Intelligence – Methodology. 2. Logic, symbolic and
Mathematical. 3. Reasoning. I. Title. II. Series.
0335.R35 1988 006.3 88-9524

ISBN 978-0-521-35236-9 hardback
ISBN 978-0-521-42421-9 paperback

CONTENTS

PREFACE

Over the last few years, more and more AI research papers have used "logical" formalisms for their knowledge representation languages. Well-known AI techniques, such as the use of "frames" and defaults, have been treated with the tools of logic. At the same time, techniques from formal logic have been borrowed by AI researchers for their own ends, often in applications far removed from the context in which they were originally developed.

Much of this work has been very useful and constructive. As a result of the increasing emphasis on formal analysis of languages for knowledge representation, we now have a much better feel for the tasks that particular languages are good (or bad) for. We are also more aware of the problems that can arise if you just invent notations without ensuring that they will behave the way you want. It has become apparent that the concerns that motivated the development of formal logic in mathematics are just as important for AI. We need to know that our knowledge representation languages are well-behaved, and we need to know about their computational properties.

At the same time, it looks as though some of the papers using such formalisms are merely disguising the poverty or unoriginality of the work being reported. It seems as though you can make your program respectable if you describe it using a dense logical notation, even if it doesn't actually do anything interesting.

The prevalence of both sorts of logic-based work in AI causes problems to the newcomer to the field, and even to experienced practitioners who don't happen to have a strong background in mathematical logic. How can you work out whether someone is using a notation to good effect, or whether they are just trying to blind you? And if they are using it for a good reason, what does it mean anyway? The aim of this book is to bring together the major *useful* applications of logic in

AI in one place, with as consistent a notation as possible and with no prejudices for or against particular schools of thought. Since I am not, in this book, arguing for any specific theory which I have invested a great deal of time and effort in, I hope I can report objectively on the most widely used theories. It is almost certain that some of my own views will creep in, particularly in the specific framework I shall be using to try to unify the various theories, but at least I am not currently arguing for one language at the expense of another. My intention is to enable my readers to conduct the arguments for themselves, not to tell them the answer.

The material is fairly dense. It is also organised in a way that reflects its origins in a two-term course for graduate AI students. The book starts with a broad overview of what logic is about, and of what it is that makes a language a logic. The next two chapters cover classical logic, with comparatively little discussion of its relevance to AI. The aim of these two chapters is to establish the specific notation used in the remainder of the book, and to present the major meta-results about soundness, completeness and decidability. Some of the proofs of meta-theorems in this part of book are rather intricate. It is quite possible to read the rest of the book despite skipping these proofs, particularly the ones towards the end of Chapter 3. Of course if you do decide to skip them, you won't *really* be in a position to understand some of the debate about the merits or otherwise of particular languages; but it won't mean that you can't follow the rest of the book at all. Chapter 4 deals with theorem proving for classical logic. It is just about possible to read Chapter 4 without reading any of what comes before or after it; and to read what comes before and after without reading Chapter 4. Chapters 5, 6, 7 and 8 discuss ways in which classical logic is inadequate for describing the phenomena which AI systems have to deal with. The topics covered in these chapters are all concerned with ways of *extending* classical logic so that it is possible to describe new phenomena, or make new distinctions. These chapters do build on one another, and it would not be a good idea to try to read them in isolation, or in some other order. Finally Chapter 9 deals with two formalisms which are predicated on the assumption that classical logic is just plain *wrong*.

Acknowledgements
In writing this book, I have had a great deal of help from a lot of people. I am particularly grateful to everyone to whom I have ever taught logic, especially Helen Chappel and Ho Mun Chan, on whom I tried out the first few drafts; to Ros Barrett, for checking the text and improving my presentations; and to Fairouz Kamereddine for her comments on the penultimate draft.

Allan Ramsay, Brighton, 1988.

Note: the book contains a lot of notation, much of which will be new to most readers. All of it is introduced as it is required. All of it, that is, apart from the symbol ∎. This is a punctuation mark indicating the end of a proof.