Representations of Commonsense Knowledge

Ernest Davis
Courant Institute for Mathematical Sciences

Morgan Kaufmann Publishers, Inc.
San Mateo, California

Sponsoring Editor *Michael B. Morgan*
Production Editor *Sharon E. Montooth*
Cover Designer *Joy Dickenson*
Copyeditor *Linda Medoff*
Compositor/Artist *Technically Speaking Publications*

Credits:

Figure 6.17: Dana H. Ballard/Christopher M. Brown, *Computer Vision*, copyright 1982, p. 250. Reprinted by permission of Prentice-Hall, Inc., Englewood Cliffs, NJ.

Figure 6.20: This figure first appeared in the journal, *Artificial Intelligence* published by North-Holland Publishing Co., Amsterdam. © 1981 by North-Holland Publishing Co. It is reprinted here by permission of the publisher.

Figure 7.8: This figure first appeared in *Artificial Intelligence In Engineering Journal*, **Computational Mechanics Publications**. The article in which it appeared is "A Logical Framework for Commonsense Predictions of Solid Object Behavior," by Ernest Davis, pp. 125-140.

Morgan Kaufmann Publishers, Inc.
Editorial Office:
2929 Campus Drive, Suite 260
San Mateo, CA 94403

94 93 92 91 5 4 3 2

Library of Congress Cataloging-in-Publication Data
Davis, Ernest.
Representations of commonsense knowledge / Ernest Davis.
p. cm. -- (Morgan Kaufmann series in representation and reasoning, ISSN 1046-9567)
Includes bibliographical references and index.
ISBN 1-55860-033-7
1. Artificial intelligence. 2. Common sense. 3. Reasoning.
I. Title. II. Series.
Q335.D37 1990
006.3--dc20

90-41289
CIP

Dedication

To my parents
Philip and Hadassah Davis

Preface

A major problem in artificial intelligence is to endow computers with commonsense knowledge of the world and with the ability to use that knowledge sensibly. A large body of research has studied this problem through careful analysis of typical examples of reasoning in a variety of commonsense domains. The immediate aim of this research is to develop a rich language for expressing commonsense knowledge, and inference techniques for carrying out commonsense reasoning. This book provides an introduction and a survey of this body of research. It is, to the best of my knowledge, the first book to attempt this.

The book is designed to be used as a textbook for a one-semester graduate course on knowledge representation. (Drafts of chapters have been used for courses at New York University, Brown, and Yale.) The one absolute prerequisite is that the student be familiar with the notation and the meaning of the predicate calculus (first-order logic). A review of the predicate calculus is given in Section 2.3, but too briefly to be useful to the student who is not already familiar with it. It is not necessary that the student be familiar with either metalogic (topics such as soundness and completeness) or with computational logic (topics such as Skolemization and resolution). It will be very helpful if the student has had a general course in AI, though more for general motivation than for specific content. Mathematical sophistication is also an asset, particularly in reading Chapters 3, 4, and 6.

The following diagram shows the interdependence of chapters:

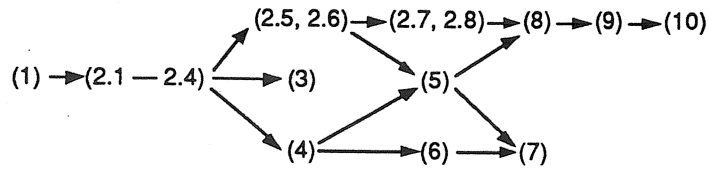

Additional dependencies of individual sections are indicated in the text. These dependencies should not be taken too seriously; in particular, the reader who finds Chapter 2 heavy going should not, on that account, be discouraged from reading the rest of the book.

Exercises are provided at the end of each chapter. Difficult exercises are marked with an asterisk. Instructors assigning starred exercises should keep in mind that they vary quite widely in difficulty, length, and the degree of prior knowledge, particularly mathematical knowledge, that they require. Some of the exercises contain results of

moderate interest that are not mentioned elsewhere in the text; the reader, therefore, may find it worth his while to glance through them even if he has no intention of working them out.

When I began work on this book in early 1985, the subject of domain-specific representations was somewhat obscure. The existing textbooks barely treated the issue, and almost no significant collections of papers had been published. The student or teacher was therefore obliged to search through journals, conference proceedings, technical reports, and unpublished papers to collect the important work in the area; and he had to form his own synthesis of the many different outlooks and techniques. The field was also small. When I planned this book, I thought I could survey virtually the entire relevant AI literature.

All this has changed spectacularly in the last five years. New textbooks, particularly [Charniak and McDermott 1985] and [Genesereth and Nilsson 1987] treat domain-specific representations at considerable length. Numerous collections of research papers, both on knowledge representation generally and on various subareas, have been published in book form, greatly simplifying the student's literature search.

At the same time, research in the area has expanded explosively. More than half the AI papers cited in the bibliography of this book were first published in 1985 or later. Today, a comprehensive bibliography would be a very substantial undertaking and a comprehensive survey would be nearly impossible.

Nonetheless, I feel that this book serves important functions. The student who has completed this book will be able to read any but the most narrowly technical paper published in the area. The researcher studying commonsense representations can find here a rich vocabulary of primitives presented within an integrated framework, and a collection of domain axioms and techniques that supports a variety of nontrivial commonsense inferences. The integration achieved here of theories of various commonsense domains makes it possible to get an overall view of how much commonsense reasoning can currently be expressed in computational terms. Compared to the powers of human commonsense reasoning, of course, we are just scratching the surface, but I think that we have made enough progress to be optimistic.

Three important omissions should be noted:

1. Domain-independent architectures for knowledge representations: semantic nets, production systems, logic programming, and so on.
2. Representations of knowledge based on linguistic considerations.

3. "Scruffy" representation of human interactions, particularly those developed by Roger Schank and his students.

The first two categories are omitted for the usual reasons: lack of time, energy, and knowledge, and a personal judgment that these issues were not central to the purposes of this book. In the case of the first category, I also felt that the subject was well covered in other texts. This is certainly not the case (as far as I know) for linguistically derived representations; a systematic survey of these would be of great value, but it is outside my personal competence.

The omission of Schankian representations is more serious. It is, in fact, a substantial disappointment to me; one of my original purposes in this book was to show how "scruffy" representation could be incorporated into a "neat" theory. In the event, however, I found this integration very hard to achieve. The problem is not so much technical; it is not difficult, using Procrustean methods, to squeeze Schankian primitives into neat categories. Rather, it is a problem of conflicting criteria in choosing primitives and inference rules. I still feel that this integration is one of the major problems in knowledge representation, but I have not found any satisfactory solution. (Section 9.5 gives a short account of Schank's categorization of goals; Appendix 10.A gives a summary of conceptual dependency.)

My friends and colleagues have been extraordinarily generous in helping me with suggestions, criticisms, and encouragement. Above all, I thank Drew McDermott, who taught me most of what I know about AI; and Leora Morgenstern, who read every draft of every chapter, and whose suggestions and comments pervade the text. Thanks also to

Sanjaya Addanki	S. Bhasker	Ron Brachman
Pasquale Caianiello	Eugene Charniak	Joey Davis
Philip Davis	Tom Dean	Sam Fleischacker
Hector Geffner	Benjamin Grosof	Pat Hayes
Jerry Hobbs	Leo Joskowicz	Ken Klein
Tomas Ksiezyk	Larry Manevitz	Chris Riesbeck
Yoav Shoham	John Sterling	Arthur Walker
Dan Weld		

The writing of this book has been supported in part by NSF grants numbers DCR-8402309, DCR-8603758, and IRI-8801529.

Finally, I thank my family for their support and encouragement. Most of all, thanks to Bianca, for her patience and love.

Contents

List of Tables

List of Figures

List of Named Axioms