# Approaches to Knowledge Representation

*An Introduction*


*Edited by*
G. A. Ringland
*and* D. A. Duce

# Approaches to Knowledge Representation: An Introduction

*Edited by*
**G. A. Ringland**
*and*
**D. A. Duce**
*Rutherford Appleton Laboratory, England*

**RESEARCH STUDIES PRESS LTD.**
Letchworth, Hertfordshire, England

**JOHN WILEY & SONS INC.**
New York · Chichester · Toronto · Brisbane · Singapore

RESEARCH STUDIES PRESS LTD.
58B Station Road, Letchworth, Herts. SG6 3BE, England

Copyright © 1988, by Research Studies Press Ltd.

All rights reserved.

No part of this book may be reproduced by any means,
nor transmitted, nor translated into a machine language
without the written permission of the publisher.

**Marketing and Distribution:**

*Australia, New Zealand, South-east Asia:*
Jacaranda-Wiley Ltd., Jacaranda Press
JOHN WILEY & SONS INC.
GPO Box 859, Brisbane, Queensland 4001, Australia

*Canada:*
JOHN WILEY & SONS CANADA LIMITED
22 Worcester Road, Rexdale, Ontario, Canada

*Europe, Africa:*
JOHN WILEY & SONS LIMITED
Baffins Lane, Chichester, West Sussex, England

*North and South America and the rest of the world:*
JOHN WILEY & SONS INC.
605 Third Avenue, New York, NY 10158, USA

**Library of Congress Cataloging in Publication Data**

Approaches to knowledge representation: an introduction/edited by
G. A. Ringland and D. A. Duce.
p. cm. — (Knowledge-based and expert systems series; 1)
Bibliography: p.
Includes index.
ISBN 0 471 91785 0 (Wiley)
1. Expert systems (Computer science) 2. Artificial intelligence.
I. Ringland, G. A. (Gordon A.) II. Duce, David A. III. Series.
QA76.76.E95A669 1988
006.3'3—dc19 87-32331

**British Library Cataloguing in Publication Data**

Approaches to knowledge representation: an
introduction.—(Knowledge-based and
expert systems series).
1. Expert systems (Computer science)
I. Ringland, G.A. II. Duce, David A.
III. Series
006.3'3 QA76.76.E95

ISBN 0 86380 064 5
ISBN 0 471 91785 0 Wiley

ISBN 0 86380 064 5 (Research Studies Press Ltd.)
ISBN 0 471 91785 0 (John Wiley & Sons Inc.)

Printed in Great Britain by SRP Ltd., Exeter

# Preface

Knowledge Representation is the keystone of the Artificial Intelligence enterprise, and systems utilizing AI techniques. Any project with a knowledge based content must choose some way of representing that knowledge, yet too rarely is this choice informed or even conscious.

This book originated from a series of lectures on Knowledge Representation given by the authors at Rutherford Appleton Laboratory. The aim is to explain and analyse a wide range of approaches to Knowledge Representation to assist in the process of rational design for knowledge based systems. The book is divided into three parts.

- The first is a discussion of the standard approaches to knowledge representation: logic, semantic networks, frames and rule based systems.
- The second is a discussion of how we, as humans, appear to represent knowledge.
- Finally a selection of more advanced topics is presented - the representation of time, meta-knowledge, conceptual graphs, issues of computational tractability, and functional approaches.

The intended audience is final year undergraduates, first year graduate students and computer professionals who are beginning to work in the areas of Knowledge Engineering and Artificial Intelligence.

**Acknowledgements**

Special thanks are due to Bob Hopgood for reading and commenting on the first draft of this volume and to Ronald Brachman of AT&T for his very perceptive comments on parts of the first draft.

We are particularly indebted to Amanda Hopgood and Charlie Kwong for producing a large number of the diagrams, to Alex Goodall of ESI Ltd. (Series Editor) and Veronica Wallace of Research Studies Press for their patience and help during the preparation of this book.

**Credits**

Chapter 3. Figure 1 adopted from M.R. Quillian "Semantic Memory" in M. Minsky (Editor) *Semantic Information Processing* (1968) with permission from MIT Press. Figure 2 reproduced with permission from McGraw-Hill Book Company from P.H. Winston *The Psychology of Computer Vision* (1975). Figure 3 adapted with permission from *Computer Models of Thought and Language* by R.C. Schank and K.M. Colby, copyright 1973 W.H. Freeman and Company. Figure 7 reproduced with permission from Elsevier Science Publishers B.V. from L.K. Schubert "Extending the expressive power of semantic networks", *Artificial Intelligence* 7(2), pp.163-198, 1976. Figure 8 is used by permission of the International Joint Conference on Artificial Intelligence, Inc. from P.J. Hayes, "On Semantic Nets, Frames and Associations", *Proc. 5th IJCAI* 1977. Copies of the Proceedings are available from Morgan Kaufmann Publishers, Inc., 95 First Street, Los Altos, CA 94022, U.S.A. Figure 9 is reproduced with permission from R.J. Brachman and J.G. Schmolze "An overview of the KL-ONE knowledge representation scheme", *Cognitive Science* **9**(2), pp.171-216, 1985.

Chapter 6. Figure 1 is reproduced with permission from the American Association for the Advancement of Science from R.N. Shepard and J. Metzler "Mental rotation of three-dimensional objects", *Science* **171**, pp.701-3, 1971. Figures 2 and 4 are reproduced with permission from Academic Press from L.A. Cooper and R.N. Shepard in *Visual Information Processing*, ed. W.G. Chase (1973) and R.N. Shepard and C. Feng "A chronometric study of mental paper folding", *Cognitive Psychology*, **3**, pp.228-243, 1972, respectively. Figure 5 is reproduced with permission from L.R. Brooks "Spatial and verbal components of the act of recall", *Canadian J. Psychology*, **22**, pp.349-368, 1968.

Chapter 7. Figures 10 and 11 are reproduced from J.F. Sowa, *Conceptual Structures: Information Processing in Mind and Machine* (1983), with permission from Addison-Wesley Publishing Co. Inc.

Chapter 10. Figure 4 is reproduced with permission of the Institute of Electrical and Electronic Engineers from J. Mylopoulos, T. Shibahara and J.K. Tsotsos, "Building Knowledge-based Systems: the PSN experience", *IEEE Computer*, **16**(10), copyright © IEEE 1983.

Gordon A. Ringland
David A. Duce

14 October 1987

试读结束，需要全本PDF请购买 www.ertongbook.com

# Table of Contents




All authors are located at the Rutherford Appleton Laboratory, Chilton, Didcot, OXON OX11 0QX, U.K.

# 1 Background and Introduction

*David Duce and Gordon Ringland*

## 1.1 Background

There is a sense in which every computer program contains knowledge about the problem it is solving. A program for solving differential equations, for example, certainly contains knowledge about that particular problem domain. The knowledge is in the particular algorithms the program employs and the decision procedure which determines which algorithm to employ in a particular set of circumstances. However, it is a characteristic of most computer programs that the knowledge they contain is not represented explicitly and cannot be readily expanded or manipulated. Knowledge is in a sense projected onto the program, like a 3-Dimensional image being projected onto a 2-Dimensional surface, and cannot be reconstructed. Given a "traditional " payroll program it would be only possible to make fragmentary deductions about, say, statutory sick pay legislation, yet this is a part of the knowledge on which the program is based and which was used in the construction of the program.

This scenario is to be contrasted with the field of Artificial Intelligence (AI) where the concern is to "write down descriptions of the world in such a way that an intelligent machine can come to new conclusions about its environment by formally manipulating these descriptions" (Brachman and Levesque, 1985a). As Sloman (1979) remarks, "work in Artificial Intelligence, whether aimed at modelling human minds or designing smart machines, necessarily includes a study of knowledge. General knowledge about how knowledge is acquired, represented and used, has to be embodied in flexible systems which can be extended, or which can explain their actions. A machine which communicates effectively with a variety of humans will have to use information about what people can be expected to know in various circumstances".

Jackson (1986) in his excellent book *Introduction to Expert Systems* gives a very succinct overview of AI. He identifies three periods in the development of AI, the Classical Period, the Romantic Period and the Modern Period. He identifies the Classical Period with the game playing and theorem proving programs that were written soon after the advent of digital computers. The game playing (for example, chess) programs of this era were based on the notion of searching a state space. Problems were formulated in terms of a starting state (e.g. the initial state of a chess board), a test for detecting final states or solutions (e.g. the rules for checkmate in chess), and a set of operations that can be applied to change the current state (for example, the legal moves in chess). In any but the simplest of cases, an exhaustive search of the state spaces was infeasible and the trick then was to find some means of guiding the search. This led to the use of rules of thumb or heuristics, that could be used to guide the search in specific domains. Chess-playing programs constructed according to this paradigm cannot be said to explicitly represent the knowledge the chessmaster has about the game and the strategies he uses to reason about this knowledge.

Similar considerations apply to theorem proving systems of this era. Jackson describes the most important discoveries of this period as the twin realizations that (a) problems of whatever kind could, in principle, be reduced to search problems providing that they could be formalized in terms of a starting state, an end state and a set of operations for generating new states, but (b) that the search had to be guided by some representation of knowledge about the domain of the problem. In most cases it was felt necessary to have some explicit representation of knowledge about the objects, properties and actions associated with the domain or to have a global problem solving strategy.

The Romantic Period is identified with the research in computer understanding that went on between the mid-1960's and mid-1970's. Whatever beliefs one may hold about the possibility of a computer understanding anything, the ability to represent knowledge about real or imaginary worlds and reason using these representations is certainly a prerequisite for understanding. Much research was devoted in this period to the development of general frameworks for encoding both specific facts and general principles about the world, and although the whole enterprise turned out to be a very nontrivial exercise, many of the approaches to knowledge representation to be described in this book have their origins in this period.

The Modern Period covers the latter half of the 1970's to the present day. There has been a growing conviction that the power of a problem solver lies in the explicit representation of knowledge that the program can access, rather than in a sophisticated mechanism for drawing inferences from the knowledge. This period has seen the development of a number of expert systems which perform well on non-trivial tasks. These programs generally have two components, a *knowledge base* which contains the representation of domain specific knowledge, and an *inference engine* which performs the reasoning. Jackson observes that these systems tend to work best in areas where there is a substantial body of knowledge connecting situations to actions. Deeper representations of the domain in terms of spatial, causal or temporal models are avoided, but these are problems that a general knowledge representation system cannot side-step quite so easily.

## 1.2 The Knowledge Representation Problem

Brachman and Levesque in their introduction to *Readings in Knowledge Representation* (1985a) remark that the notion of knowledge representation is essentially an easy one to understand. It simply has to do with writing down, in some language or communications medium, descriptions or pictures that correspond in some salient way to the world or a state of the world. As in other areas of computer science, it is also necessary to consider the ways in which the representation is to be manipulated and the uses to which it is to be put. As remarked earlier, the primary reason for wanting to represent knowledge is so that a machine can come to new conclusions about its environment by manipulating the representation.

The first ingredient of the knowledge representation problem is to find a *knowledge representation language*, that is some formal language in which domains of knowledge can be described. Most systems of practical interest then need to be able to provide their users with access to the facts implicit in the knowledge base as well as those stored explicitly, and thus it is necessary to have a component of the knowledge representation that can perform automatic *inferences* for the user. The third component of the knowledge representation problem is how to capture the detailed knowledge base that represents the system's understanding of its domain. This latter problem is beyond the scope of this book, however.

David Israel characterized the knowledge representation problem as follows:

> All parties to the debate agree that a central goal of research is that computers must somehow come to "know" a good deal of what every human being knows about the world and about the organisms, natural or artificial, that inhabit it. This body of knowledge - indefinite no doubt, in its boundaries - goes by the name "common-sense". The

problem we face is how to impart such knowledge to a robot. That is, how do we design a robot with a reasoning capacity sufficiently powerful and fruitful that when provided with some subbody of this knowledge, the robot will be able to generate enough of the rest to intelligently adapt to and exploit its environment? We can assume that most, if not all, common-sense knowledge is general, as is the knowledge that objects fall unless they are supported, that physical objects do not suddenly disappear, and that one can get wet in the rain.

The following simple example, given by Minsky, points out that knowledge representation is not a simple problem:

> The only time when you can say something like, "if *a* and *b* are integers, then *a* plus *b* always equals *b* plus *a*", is in mathematics. Consider a fact like "Birds can fly". If you think that common-sense reasoning is like logical reasoning, then you believe there are general principles that state, "If Joe is a bird and birds can fly, then Joe can fly". Suppose Joe is an ostrich or penguin? Well we can axiomatize and say if Joe is a bird and Joe is not an ostrich or a penguin, Joe can fly. But suppose Joe is dead? Or suppose Joe has his feet set in concrete?

It is worth exploring this theme a little further. Some domains of knowledge, for example mathematical knowledge, are well-behaved in a certain sense, and are relatively straightforward to deal with. For example, a triangle is a 3-sided polygon, or the sum of the interior angles of a triangle is 180°. These facts are true of all triangles and can be used as definitions of the concept of a triangle.

For other domains of knowledge, it is not quite so straightforward. Some concepts, for example *bachelor*, have an explicit definition "a man who has never married" (at least that is true when the terms are used strictly!). However, the majority of names do not have simple definitions of this form. An important class of objects are *natural kinds* (naturally occurring species), for example *lemon*, and *elephant*. The book *Naming, Necessity and Natural Kinds* (Schwartz, 1977) contains a fascinating collection of papers on this subject which is well worth studying, if only to remind oneself that the problems of knowledge representation did not arise with the advent of digital computers, but have long been studied by philosophers whose writings ought not to be ignored by computer scientists.

Putnam in his paper "Is Semantics Possible?" in the above volume, looks in detail at natural kind objects. In the traditional philosophical view, the meaning of, say, "lemon", is given by specifying a conjunction of *properties*, akin to the definition of triangle. A lemon is something that has all of the properties in the definition. Putnam and the other authors in (Schwartz, 1977) challenge this traditional view. Suppose the defining characteristics of a lemon are "colour lemon", "tart taste" etc. The problem is that a natural kind may have abnormal members, for example there are green fruits that everyone would agree are lemons, and elephants with three legs are still elephants. It is argued that nouns meant to designate natural kinds do not have their extensions (the set of things to which they refer) determined by a finite number of concepts.

Suffice it to say in this chapter, that it is important when choosing a knowledge representation scheme for a particular domain of knowledge, to consider the types of objects in the domain.

Some of the issues that arise in knowledge representation are summarized below to give more of a feeling for the problems.

(1) *Expressive adequacy.* Is a particular knowledge representation scheme sufficiently powerful? What knowledge can and cannot particular schemes represent?

(2) *Reasoning efficiency.* Like all representation problems in computer science, a scheme that represents all knowledge of interest and is sufficient to allow any fact of interest to be inferred by no means guarantees that it will be possible to perform the inference in an acceptable time. There is generally a tradeoff between expressive adequacy and reasoning efficiency.

(3) *Primitives.* What are the primitives (if any) in knowledge representation? What primitives should be provided in a system and at what level?

(4) *Meta-representation.* How do we structure the knowledge in a knowledge base and how do we represent knowledge about this structure in the knowledge base?

(5) *Incompleteness.* What can be left unsaid about a domain and how do you perform inferencing over incomplete knowledge and revise earlier inferences in the light of later, more complete, knowledge?

(6) *Real-world knowledge.* How can we deal with attitudes such as beliefs, desires and intentions? How do we avoid the paradoxes that accompany self-referential propositions?

The remainder of the first part of this book describes four approaches to the knowledge representation problem which have acquired some degree of acceptability amongst researchers in the field. The four approaches are: logic, semantic nets, frames, logic and rule based systems. Subsequent chapters deal with each of these approaches in turn. The second part of the book covers some current research directions, and problems common to all of these basic approaches, for example the representation of time and the trade-off between expressive power and the computational efficiency of

试读结束，需要全本PDF请购买 www.ertongbook.com

inferencing.

The next section gives a brief introduction to the four basic approaches.

## 1.3 Overview of the Basic Approaches

### 1.3.1 Logic

Mathematical logic is an attempt to make rigorous the reasoning process involved in mathematics. The starting point is the introduction of a symbolic language whose symbols have precisely stated meanings and uses. The next step is to define the rules by which these symbols can be combined and manipulated and then the properties of the resulting formal system are explored. Chapter 2 gives a detailed introduction to various systems of mathematical logic and their application to knowledge representation. In this introduction, we will give a flavour for the approach in a very informal style.

A recent paper by Sergot *et al.* (1986) describes the use of a certain system of logic to describe a large part of the British Nationality Act, 1981. The system of logic is known as definite Horn Clauses, which are essentially rules of the form:

$$A \text{ if } B_1 \text{ and } B_2 \text{ and } \cdots B_n$$

which have exactly one conclusion $A$, but zero or more conditions $B$. A simple example of a Horn clause is the following:

(Socrates is mortal) if (Socrates is a man)

The first clause of the British Nationality Act is as follows:

1.-(1).A person born in the United Kingdom after commencement shall be a British citizen if at the time of birth his father or mother is
(a) a British citizen; or
(b) settled in the United Kingdom.

Clause 1.-(1)(a) is represented as a first approximation by:

($x$ is a British citizen)
if ($x$ was born in the U.K.)
and ($x$ was born on date $y$)
and ($y$ is after or on commencement)
and ($z$ is a parent of $x$)
and ($z$ is a British citizen on date $y$)

The symbols $x$, $y$ and $z$ are variables.

Using a slight extension of this mathematical apparatus, a major part of the British Nationality Act was represented. Having obtained such a representation, it can then be manipulated using the rules of logical inference, appropriate to this system of logic, so that answers to queries such as "is Peter a British citizen on 16 January 1984 given that he was born on 3 May 1983 in the U.K. and is still alive and his father William ...", can be given.

### 1.3.2 Semantic Networks

The study of semantics is an attempt to describe the concepts behind word meanings and the ways in which such meanings interact. It is such a description which semantic networks were designed to provide. A network is a net or graph of nodes joined by links. The nodes in a semantic network usually represent concepts or meanings (e.g. BOOK, GREEN) and the links (or labelled directed arcs) usually represent relations (e.g., a book IS COLOURED green).

Semantic networks may be loosely related to predicate calculus by the following substitution: *terms* are replaced by *nodes* and *relations* by *labelled directed arcs*.

A large number of semantic networks have been developed as variations on this simple pattern since Quillian (1968) first used one in a computer system. These networks share few assumptions, although they nearly all represent the relations between concepts using a semantic representation consisting of a network of links between nodes, a set of interpretative processes that operate on the network, and a parser. They also show a general commitment to parsimony.

The most often used link in semantic networks was introduced in Quillian's system to show that one concept is an example of another (e.g. canary IS-A bird). More recent systems have chosen their link and node types on the basis of epistemelogical concerns about how the knowledge will be used. These have shown that even the apparently simple IS-A relationship is more complex than had been previously believed.

Recent developments in semantic networks together with work on the theoretical underpinnings of this approach are reviewed in the chapter by Mac Randal.

### 1.3.3 Frames

The use of nodes and links to represent concepts and relations seems straightforward, but contains many pitfalls.

Some designers of network systems were not too careful about the way in which they assigned meanings to nodes. Thus, a type node labelled "elephant" might well stand for the concept of elephant, the class of all elephants, or a typical elephant. Similarly token nodes labelled elephant were open to interpretation as a particular elephant, an arbitrary elephant etc. Different interpretations support different sets of inferences and so the distinctions are important. There was thus a sense in which semantic network formalisms were logically inadequate in that they could not make many of the distinctions that can be easily made in mathematical logic, for example between a particular elephant, all elephants, no elephant etc.

Frames are ways of grouping information in terms of a record of "slots" and "fillers". The record can be thought of as a node in a network, with a special slot filled by the name of the object that the node stands for and the other slots filled with the values of various common attributes associated with such an object. Frames are particularly useful when used to represent knowledge of certain stereotypical concepts or events. The intuition here is that the human brain is less concerned with defining strictly the properties that entities must have in order to be considered as exemplars of some category, and more concerned with the salient properties associated with objects that are typical of their class.

Frame systems reason about classes of objects by using stereotypical representations of knowledge which usually will have to be modified in some way to capture the complexities of the real world, for example that birds can fly, but emus cannot. The idea here is that the properties in the higher levels of the system are fixed, but the lower levels can inherit values from higher up the hierarchy or can be filled with specific values if the "default" fillers are known to be inappropriate.

### 1.3.4 Rule Based Systems

A classic way to represent human knowledge is the use of IF/THEN rules. The satisfaction of the rule antecedents gives rise to the execution of the consequents - some action is performed. Such production rule systems have been successfully used to model human problem-solving activity and adaptive behaviour.

More recently, substantial knowledge-based systems have been constructed using this formalism, for example the R1/XCON computer configuration system, implemented in the OPS5 production rule language. Chapter 5 describes the basic operations of a production system and the problems which arise in systems involving large numbers of rules, as well as considering the suitability of this formalism as a general knowledge representation.

## 1.4 Psychological Studies of Knowledge Representation

The second part of the book is a review of how we, as humans, appear to represent knowledge.

In this chapter the schemes which have been suggested as being those used to represent knowledge in human memory are reviewed. These include the use of frames, schema, semantic nets and production rules described in the earlier chapters. Instantiations of these are described for which both the representations and processes acting on them are specified in sufficient detail to enable experimentally testable hypotheses to be drawn. Experimental evidence is presented which supports an argument that schemes using only one of these representation mechanisms are inadequate to account for the full range of phenomena exhibited in human performance, although individual models can account for the specific sets of phenomena which they are intended to address.

A class of analogical representations is introduced which has not been described in earlier chapters but which are capable of supporting the phenomenon of visual imagery. Evidence is presented as to the use of imagery by humans and the nature of the representations which would have to support it. This suggests that although it is possible to account for visual imagery by processes acting on a propositional representation, it seems more likely that some form of analogical representation is used by humans.

One use of analogical representations is to form models of situations so that reasoning can be performed on them. Johnson-Laird's (1983) suggestions as to how such *mental models* could be used to support inference are described, along with findings which suggest the use of both propositional and such analogical representations by humans. As well as describing the limitations of suggested representation schemes and providing evidence that supports the use of multiple forms of representation, this chapter provides a set of phenomena for which any representation scheme will have to account if it is to address the range of human performance.

## 1.5 More Advanced Topics in Knowledge Representation

The third part of the book reviews a selection of more advanced topics.

### 1.5.1 Conceptual Graphs

In Chapter 7 Jackman and Pavelin give an overview of the basic concepts of the conceptual graph knowledge representation language. This includes the concept of the conceptual graph, the type hierarchy, the basic operations that may be performed on conceptual graphs, and logical deduction. Reference is also made to the "maximal join" - one of the fundamental derived operations in the language. This operation would appear to be equivalent to the graph equivalent (with a type hierarchy) of unification.

### 1.5.2 The Explicit Representation of Control Knowledge

Production systems have been used in many knowledge-based systems to model human expertise in classification. For example, the MYCIN family of expert systems can identify which microbial organisms are producing symptoms of disease in a patient. Important criticisms of such systems have been made by Clancey and others. Although the systems effectively "do the job" of the expert physician, much of the knowledge has been compiled, which is to say that it has been compressed and restructured into effective procedures. Bainbridge in Chapter 8 shows this makes it difficult to re-use the knowledge in explanation and knowledge acquisition subsystems, since the knowledge is implicit and therefore unavailable.

An important research area involves reconstructing these systems to make the knowledge explicit and available for use, and from these implementations extracting general principles for making better expert systems which more effectively represent the knowledge in their domain.

### 1.5.3 Representing Time

One of the most fundamental, and deceptively simple, representations that humans have is that of time. A great deal of effort has been expended on attempting to formulate temporal representations for use by knowledge-based systems. Chapter 9 considers first the basic issues in the representation of time, such as the choice of point or interval representations, the treatment of fuzziness and granularity, and the problem of persistence. A number of approaches are then presented, with reference to the systems in which they have been used or the contexts in which they are appropriate. State-space modelling, date-based methods and before/after chains are all covered, along with temporal logics, which have attempted to place representations of time on a formal foundation.

### 1.5.4 Functional Approaches

An important approach to knowledge representation is the functional approach pioneered by Levesque and Brachman. There is a relation to mainstream computer science in that a knowledge base is regarded almost as an abstract data type with a set of operations defining the services it provides. The approach is motivated by the misuses or misinterpretations of knowledge representation formalisms which can occur when the user is allowed unrestricted access to representational structures: for instance, the nodes and links of a semantic net. Chapter 10 discusses some early work, and then describes Levesque's formalization in which he defines operators TELL and ASK for interacting with a knowledge representation system. Finally, the KRYPTON system is dealt with. It is the most advanced implementation of functional ideas, and it also incorporates multiple representations in having a taxonomic component for defining absolute relationships and an assertional component for making statements.

### 1.5.5 Expressive Power and Computability

There is a fundamental difference between a knowledge representation system and a database: the former will in general perform inferencing of some kind in order to answer queries about what is represented, while the latter is limited to retrieving the facts it contains. Databases cannot therefore represent incomplete information, for everything must be stored explicitly. Knowledge representation systems are more expressive, and their inference capabilities mean that they can act on incomplete knowledge. Indeed, when there is incomplete knowledge, queries to a database concern no more than what the database happens to contain; only a knowledge representation system can go further and attempt to deal with the world it represents. Of course the price to pay is in the computational effort needed to answer queries - the trade-off between the two factors is discussed in Chapter 11 by Williams and Lambert.

It is well-known that full first-order logic is not decidable, that is, a theorem prover cannot be guaranteed to terminate. Restricting the expressive power of the representation language results in systems that exhibit various degrees of tractability: though decidable, some are NP-complete, while others, less expressive, admit inferencing algorithms that operate in polynomial time. A number of the knowledge representation schemes described earlier in the book are discussed in these terms. It is not yet understood precisely how the tractability of a knowledge representation system depends on its expressiveness though there are some indications, but the trade-off may have important implications for our view of what service is expected of such systems.

试读结束，需要全本PDF请购买 www.ertongbook.com