
KNOWLEDGE ENGINEERING

Volume I FUNDAMENTALS



73.0
56

KNOWLEDGE ENGINEERING

Volume I FUNDAMENTALS

Hojjat Adeli, Editor

*Department of Civil Engineering
The Ohio State University*

9350019
McGraw-Hill Publishing Company

New York St. Louis San Francisco Auckland Bogotá Caracas
Hamburg Lisbon London Madrid Mexico Milan Montreal New Delhi
Oklahoma City Paris San Juan São Paulo Singapore Sydney Tokyo Toronto

9350019

This book was set in Times Roman by the College Composition Unit
in cooperation with Ruttle, Shaw & Wetherill, Inc.

The editors were B. J. Clark and John M. Morriss;

the production supervisor was Leroy A. Young.

The cover was designed by Carla Bauer.

Project supervision was done by The Total Book.

R. R. Donnelley & Sons Company was printer and binder.

KNOWLEDGE ENGINEERING

Vol. I, Fundamentals

Copyright © 1990 by McGraw-Hill, Inc. All rights reserved. Printed in the
United States of America. Except as permitted under the United States Copyright Act
of 1976, no part of this publication may be reproduced or distributed in any form or
by any means, or stored in a data base or retrieval system, without the prior written
permission of the publisher.

1 2 3 4 5 6 7 8 9 0 DOC DOC 9 5 4 3 2 1 0

ISBN 0-07-000355-6

Library of Congress Cataloging-in-Publication Data

Knowledge engineering / [edited by] Hojjat Adeli.

p. cm.

Includes index.

Contents: v. 1. Fundamentals—v. 2. Applications.

ISBN 0-07-000355-6 (v. 1)—ISBN 0-07-000357-2 (v. 2)

1. Expert systems (Computer science) I. Adeli, Hojjat, (date).

QA76.76.E95K577 1990

006.3'3—dc20

89-8338

EDITOR'S BIOGRAPHY

Currently professor of civil engineering at The Ohio State University, Hojjat Adeli received his Ph.D. from Stanford University in 1976. He is the editor-in-chief of the international journal, *Microcomputers in Civil Engineering*, and the author or editor of nearly 200 research publications including several books in the fields of knowledge engineering and expert systems, computer-aided design, parallel processing, mathematical optimization and simulation, applied mechanics, and structural engineering. He is also the editor-in-chief of the forthcoming Marcel Dekker book series *New Generation Computing*. The first two volumes of the series, *Supercomputing in Engineering Analysis* and *Parallel Processing in Computational Mechanics* are scheduled for publication in late 1990. He is listed in twelve *Who's Who's* and biographical listings including *Who's Who in the World*, *Men of Achievement*, and *International Directory of Distinguished Leadership*.

CONTRIBUTORS

Sergio J. Alvarado received his Ph.D. in computer science from the University of California, Los Angeles, in 1989. He is currently an assistant professor in computer science and director of the Artificial Intelligence Laboratory at the University of California, Davis.

Wolfgang L. Bibel is a professor of computer science at the University of British Columbia and fellow of the Canadian Institute for Advanced Research. He received his Ph.D. from the University of Munich, West Germany, in 1968. Between 1969 and 1987 he worked at the Technical University of Munich. His nearly 100 publications range over various areas in artificial intelligence such as automated deduction, program synthesis, and knowledge representation. Dr. Bibel is section editor of the *Artificial Intelligence Journal* and the associate editor of the *Journal for Symbolic Computations*. He chaired the International Joint Conference on Artificial Intelligence in 1989.

Mark Drummond received his Ph.D. in artificial intelligence from the University of Edinburgh in 1986. He subsequently worked as a postdoctoral fellow at the University of Edinburgh's AI Applications Institute. He is currently a scientist at the NASA Ames Research Center, Moffett Field, California. His current research interests include the relationship between planning, scheduling, and execution; real-time reactive scheduling in light of execution failures; and the role of dependency information in error management. Of particular interest is the Mars Rover Sample Return Mission.

Michael G. Dyer received his Ph.D. in computer science from Yale University in 1982. He is currently an associate professor of computer science and director of the Artificial Intelligence Laboratory at the University of California, Los Angeles. He is the author of *In Depth Understanding* (MIT Press, 1982), and has authored over 60 articles in books, journals, and conference proceedings, covering numerous areas of cognitive science, including computational linguistics.

tics, legal reasoning, modeling of emotions, and artificial neural networks. Dr. Dyer is on the editorial advisory board of the journals, *Expert Systems: Research and Applications*, *Knowledge-Based Systems*, and *Connection Science*.

Erdal Elver studied informatik at the Technical University of Munich from 1981 to 1986. Since 1986 he has been with Siemens AG in Munich where he is now the project leader of the PRINCESS group and works on the development of expert system shells. His special interests are expert systems and knowledge representation.

Margot Flowers did graduate work in computer science at Yale University and is currently an adjunct assistant professor in the computer science department at the University of California, Los Angeles. She has authored over 20 publications in natural language processing, connectionist modeling, legal reasoning, mechanical device comprehension, and scientific reasoning.

Brian R. Gaines is Killam Memorial Research Professor and director of the Knowledge Science Institute at the University of Calgary. He is also professor of computer science and psychology, and director of the Software Research and Development Group at the University of Calgary. His previous positions include professor of industrial engineering at the University of Toronto, technical director and deputy chair of the Monotype Corporation, and chair of the department of electrical engineering and science at the University of Essex. He received his B.S., M.A., and Ph.D. from Trinity College, Cambridge, and is a chartered engineer, and a fellow of the Institution of Electrical Engineers, the British Computer Society, and the British Psychological Society. He is editor of the *International Journal of Man-Machine Studies* and *Future Computing Systems*, as well as the *Computers and People* and *Knowledge-Based Systems* book series. He has authored over 250 papers and authored or edited 5 books in various areas of computer and human systems. His research interests include the socioeconomic dynamics of science and technology; the nature, acquisition, and transfer of knowledge; software engineering for heterogeneous systems; and expert system applications in manufacturing, sciences, and humanities.

Forouzan Golshani received his Ph.D. from Warwick University, United Kingdom, in 1982. He joined Arizona State University as an assistant professor of computer science in 1984. He has written over 40 publications in the areas of expert systems, parallel architectures, software engineering, and database management.

Mehdi T. Harandi obtained his M.Sc. and Ph.D. from the University of Manchester, England, in 1976 and 1979, respectively. In 1981 he joined the department of computer science at the University of Illinois, Urbana-

Champaign, where he is currently an associate professor of computer science. He is the editor-in-chief of the *International Journal of Expert Systems: Research and Applications*. He was the program chair of the 1987 International Workshop of Software Specification and Design. His current research interests are in the areas of software specification and design, knowledge bases and their application to software development, and distributed expert systems.

Yves Kodratoff is presently director of research at the French National center for Scientific Research. He has done extensive research in the area of machine learning and published a book on the subject recently. He was the program chair of the 1988 European Conference on Artificial Intelligence. His current research interests are program synthesis from specifications and its application to software engineering, and various aspects of machine learning, including the coupling of numeric and symbolic techniques, developing abductive and inductive tools working in the presence of a large knowledge base, and machine discovery and analogy.

Rense Lange received his Ph.D. in psychology from the University of Illinois, Urbana-Champaign in 1981. He is now working toward the completion of a Ph.D. thesis in the area of artificial intelligence at the University of Illinois. He has published in the areas of cognitive processing, problem solving, decision making, and machine learning. His research in expert systems focuses on knowledge representation and knowledge acquisition.

Yoh-Han Pao is professor of electrical engineering and computer science and director of the Center for Automation and Intelligent Systems Research at Case Western Reserve University. He was the chair of Case Western's electrical engineering department during 1969-1977 and director of the National Science Foundation Division of Electrical, Computer, and Systems Engineering during 1978-1980. Professor Pao has received numerous awards including the U.S. Emblem of Merit for Distinguished Civilian Service in 1945 and the Outstanding Educator of America award in 1972 and 1973. He was the technical editor of the *IEEE Journal of Robotics and Industrial Automation*. He is a fellow of the IEEE and the Optical Society of America. His current interests are expert systems, pattern recognition, distributed processing, and neural nets.

Josef Schneeberger studied informatik at the Technical Universities of Munich and Darmstadt. From 1985 to 1988 he was a member and leader (for one year) of the Artificial Intelligence Group of the Technical University of Munich. He is currently on the faculty of the Technical University of Darmstadt.

Philip E. Slatter is a knowledge engineer with Telecomputing plc in Oxford, England, where he has worked on large-scale expert system applications since 1985. Between 1980 and 1985 he was an analyst/programmer with the British

Steel Corporation. His academic qualifications include a B.Sc. in psychology, an M. Phil. for research into cognitive aspects of expert systems, and a Ph.D. in applied cognitive psychology. He is the author of a recent book entitled *Building Expert Systems: Cognitive Emulation* (Ellis Horwood, 1987).

Austin Tate is director of the Artificial Intelligence Applications Institute at the University of Edinburgh. Dr. Tate's research interests relate primarily to knowledge-based planning systems, novel database architectures, and knowledge representations systems. He is a member of the steering group of the United Kingdom Knowledge-Based Systems Research Club and the coordinator for its Planning Special Interest Group. He is on the editorial board of several journals and technical book series.

Ronald R. Yager is currently director of the Machine Intelligence Institute and professor of information systems at Iona College. He received his Ph.D. from the Polytechnic Institute of New York in 1968. During 1983–1984 he served at the U.S. National Science Foundation as program director in the Information Sciences program. He is the editor-in-chief of the *International Journal of Intelligent Systems*. He also serves on the editorial board of a number of other journals. He has published close to 200 articles. His current research interests include multicriteria decision making, uncertainty management in knowledge-based systems, fuzzy set theory, neural modeling, and visual databases.

PREFACE

The first volume of *Knowledge Engineering* presents state-of-the-art reviews and tutorials on fundamental aspects of knowledge engineering. The second volume complements the first by presenting applications of applied artificial intelligence (AI). The field of applied AI and knowledge engineering is very young. Students usually must refer to numerous sources to learn the fundamentals of the subject. The two volumes attempt to present summaries of the various subjects in a single document and are oriented toward practical applications. They are suitable as primary reference books in introductory courses on applied AI and knowledge engineering.

Leading and internationally recognized researchers have contributed to these volumes. We hope this effort becomes a continuing book series with future volumes concentrating on other aspects of knowledge engineering and new applications of AI.

Hojjat Adeli
Editor

CONTENTS

Contributors	ix
Preface	xiii
1 Representation of Knowledge	
W. Bibel, J. Schneeberger, and E. Elver	1
2 Rule-Based Expert Systems	
F. Golshani	28
3 Knowledge Acquisition Systems	
B. R. Gaines	52
4 Model-Based Knowledge Acquisition	
M. T. Harandi and R. Lange	103
5 Models of Expertise in Knowledge Engineering	
P. E. Slatter	130
6 AI Planning	
M. Drummond and A. Tate	157
7 Knowledge in the Form of Patterns and Neural Network Computing	
Y. H. Pao	200
8 Machine Learning	
Y. Kodratoff	226
9 Propositional Logic	
R. R. Yager	256

**10 Natural Language Processing: Computer
Comprehension of Editorial Text**

S. J. Alvarado, M. G. Dyer, and M. Flowers

Index

CHAPTER 1

REPRESENTATION OF KNOWLEDGE

W. BIBEL
J. SCHNEEBERGER
E. ELVER

1 INTRODUCTION

Engineers of all sorts are increasingly interested in including "knowledge" among the materials from which they construct their artifacts. Whereas traditional materials such as metal, wood, or plastics are tangible, knowledge is an extremely evasive stuff. Thus the need to "materialize" or *represent* knowledge arises in the first place. This chapter introduces some of the techniques that have been developed to allow for the *representation of knowledge*, especially within a machine, and discusses the major problems in such an endeavor.

As a research area in its own right, knowledge representation evolved within the field of artificial intelligence (AI), where it continues to play a central role. This should not come as a surprise since everyone would probably agree that intelligence has a lot to do with how knowledge is being handled in the human mind and, hence, in any kind of intelligent device. From this AI perspective the area started some thirty years ago.

From a more general perspective, knowledge representation has a much longer tradition rooted mainly in philosophy, logic, and psychology. In fact, both notions "knowledge" and "representation" touch on deep philosophical issues that cannot totally be ignored even from an engineer's point of view. What is knowledge? Although our presentation naturally emphasizes the more technical aspects of knowledge representation, we will nevertheless give a brief account of some of these philosophical issues in the subsequent section.

Even in its more technical aspects, knowledge representation has turned out to be a difficult area that does not seem to lend itself to easy solutions of its

main problems in the near future. The reader must therefore not expect the presentation of a few simple recipes to be observed in building the next expert system, robot device, or what have you. While researchers in this area begin to agree on what the fundamental problems are, their proposed solutions do not converge yet. For instance, there is still a vast amount to be understood before the knowledge represented in one system can be shared with another system, something that people apparently can easily achieve.

Another major difficulty in presenting this area consists of the fact that it spans a wide spectrum of research topics. In consequence, to write a sort of tutorial about this field in a single chapter is obviously not an easy task. We have chosen to do this in the following way.

First, we have restricted the main discussion to a relatively small subarea of the entire field. In particular, we took knowledge representation in its literal sense to denote ways of just representing knowledge via some formalism or another. Usually much more is covered by this notion since it is impossible to strictly separate the issues of representing from those of processing knowledge. One influences the other in an intrinsic way. To remind the reader once again of this drastic restriction, in Section 5 we offer a few key words and pointers to the important and large remaining part of the field not covered here in detail. In addition, we encourage the reader to consult the literature cited or perhaps to even start reading a more general textbook on AI such as Charniak and McDermott (1985), to mention a recent one.

The core of this chapter consists of Sections 3 and 4. There we first survey some ten formalisms that have been used extensively in the literature and in knowledge-based systems. We use logic as a sort of canonical formalism that provides meaning to all others as well. We then point out the strong points of most of these formalisms relative to each other on the basis of rational arguments rather than personal taste. While the authors' bias toward logical formalisms will not be concealed, we hope to have collected convincing arguments for this preference.

Those readers with no taste at all for the philosophically flavored issues involved might well skip the subsequent section entirely.

2 KNOWLEDGE AND ITS REPRESENTATION

2.1 Basic Concepts

Like many other general concepts, the notion of *knowledge* is ill defined and thus has a different meaning for different people. Primarily it is a cognitive notion: what we "know" is anything that is at our disposal in the conscious mind. To some extent, knowledge can be communicated among people through some sort of medium such as spoken language. For this purpose knowledge is somehow represented in such a medium. That is, we establish a

certain relationship between the represented knowledge and the cognitive phenomenon that constitutes this knowledge.

Consider the sentence "The gripper picks up the log" uttered by someone watching a scene of logging. The observer has in mind the knowledge of this scene (depicted as the middle box in Fig. 1-1) of which the sentence is a represented form (the left box in Fig. 1-1). The interpretation of the represented knowledge is the cognitive knowledge in the mind. The figure suggests that this interpretation is a function ι mapping the sentence to the cognitive knowledge.

The cognitive knowledge in this case arises from the observation of a scene in the external or real world that is depicted in the right-hand box of Fig. 1-1. We presume that cognitive knowledge is represented somehow in the brain which allows us to think of it as a represented form of the external-world scene. Hence we have the right arrow from left to right. By transitivity this yields an interpretation of the sentence in terms of the scene in the external world. Keep in mind that such an interpretation is not possible except by way of interaction with an intermediate mind, which plays the central role, as clearly demonstrated by the figure.

Although it may appear as an unusual view, the scene in the external world may well be considered as a representation of the cognitive knowledge as well. Hence we have the arrow from right to left as well. We will not dwell on the resulting one-to-one correspondence further which obviously begs many deep scientific and philosophical questions. Also we mention, as an aside, that for human beings in a sense there might be no "real" world (and thus no represented form of knowledge either) since their experiences take place exclusively in the realm of consciousness. Yet operating under the assumption of its existence clearly is a very successful "working hypothesis".

It is in such represented form only, as in the quoted sentence above, that knowledge becomes more generally accessible and in this way adopts additionally a figurative, noncognitive meaning. Although a lot of confusion may arise if one ignores the cognitive origin of the notion, the main interest from our point of view lies in this figurative meaning. When we refer to knowledge in this noncognitive sense, we also prefer to think of the interpretation ι as a function into the external world, or rather into some other representation of it; that is, ι then abbreviates the composition of the two functions illustrated in Fig. 1-1.

There are many ways of representing knowledge (as there are many kinds of knowledge). Writing text like this one provides one method of representation. We are particularly interested in ways that use a sufficiently precise no-




FIGURE 1-1
Knowledge and its representation.

tation that it can be used in, or by, a computer program. A systematic representational method satisfying this requirement is referred to as a *scheme*. A number of such schemes are discussed in the next section. Here we must mention (the language part of) logical calculi as one type of example. A particular expression in such a scheme is called a *configuration*, such as a particular logical formula in a logical calculus. In other words, a scheme is a set of configurations.

To comply with the aforementioned requirement, we consider it intrinsic to the notion of a scheme that a computer may verify whether a particular arrangement of marks is a well-formed configuration, i.e., a definite notion of *well-formedness* is provided along with the configuration of a scheme. This requirement narrows the concept of a scheme to strictly formal ways of representation and, at least for the time being, eliminates many other ways which humans have of conveying meaning such as drawings, poems, conversational English, musical performances, and so forth. This is to say that we do not exclude altogether the possibility of formalizing also these ways in some approximation.

A second important requirement on a scheme is to provide for an associated *semantic theory* that assigns a meaning to any of its possible configurations. We want to be able to understand the represented knowledge after all. This amounts to capturing in a precise way the interpretation function ι which will be understood as a mapping into the external world, as described above.

The configurations of a scheme need to be constructed in some *medium*, for instance, as marks or symbols on a piece of paper. There are all sorts of such media. And in fact it is rather difficult to give a precise characterization of this notion and its properties, but it is less of a problem for the more practical purposes.

Further, the represented knowledge typically is not just sitting there, but first has to be imprinted into the chosen medium and may then play a number of different roles. The represented knowledge may be accessed, read, changed, linked, or what have you. In other words, a number of *operations* have to be taken into consideration that are to be applied to represented knowledge.

To summarize, the goal of a suitable representation of knowledge requires the consideration of the kind of knowledge to be represented, of the medium available for the representation, of the scheme to be used for the representation, of the meanings that the symbols used in a configuration are supposed to convey, and of the operations that are to be applied in the manipulation of the represented knowledge. This leaves us with a fairly large degree of freedom since quite a variety of choices are to be made here. Moreover, the kind of knowledge as well as the necessary operations may vary quite a bit under different circumstances. It is because of this variety that there exists no such thing as *the* knowledge representation formalism in artificial intelligence or *intellectics*.¹ Rather, there are quite a few competing ones.

¹ Our preferred name for the union of artificial intelligence and cognitive science.

2.2 Desiderata in Knowledge Representation

As mentioned in the previous section, there are a number of competing schemes in knowledge representation and hot debates on which of those should be considered the "right" one. As one might expect, at present we are *not* in a position to end these debates through some convincing argument. To get at least some sort of a measure for the evaluation and comparison of the different schemes to be discussed, we state here some of the major goals to be pursued in knowledge representation. Since people provide the best model for how to deal with represented knowledge, it is a good idea to look carefully at some of the features they realize.

Let us first consider the purpose of representing knowledge such as the earlier one expressed in the sentence about the gripper picking up the log. We would like an intelligent device to interpret this sentence in such a way so that the device could answer questions about the current location of the log and availability of the gripper for further tasks and possibly notice some danger for a person standing close by and issue a warning.

What these questions demonstrate is that a lot of additional information goes along with such a sentence. A major goal in knowledge representation is to get a machine to associate such information with a sentence, like the one we consider, when it is representing the sentence internally, an extremely hard problem indeed (known as the *qualification problem*).

But even representing the sentence itself begs a number of questions. Should the sentence be represented just as it is or in some different, possibly less readable form? The behavior of human beings indicates that our brains use the second alternative since we tend to remember the gist of a conversation long after we have forgotten the exact wording. A more technical argument lies in the observation that the sentence alone is *referentially ambiguous* if it is looked at later. "Which gripper, which log are you talking about?" one would ask. So unique names have to be associated with such objects. There are further ambiguities in natural language as well as in pictures and other external sources of information that have to be resolved while representing knowledge internally.

One of the most striking features of human knowledge processing is the fact that the communication of knowledge may change behavior even though the knowledge must not necessarily contain anything like advice for how to behave. If I learned that the road I was about to take was closed, then I would immediately change my plans even though what I had learned was a fact rather than anything of the nature of a command. Apparently people can transform such knowledge into procedural behavior. Thus they are able to accumulate knowledge in an additive way and yet use it at the same time for procedural purposes.

This is not the only way that people acquire knowledge. I learned how to ride a bike, for instance, by some mechanism that we know relatively little about. It is a fact that current expert systems technology has accomplished neither of these two possible ways of acquiring knowledge to a satisfactory de-

gree, probably one of the major reasons for its relatively moderate commercial success thus far. Given their apparent relevance in human knowledge processing, I would put both *additivity* and *learning* of procedural behavior high on the list of desiderata for systems that represent and process knowledge. The scheme used for the representation should support either of these properties if possible.

One reason to feed knowledge into a computer is to aid human experts in their work. This aid will be the more effective as it becomes easier for the expert to quickly grasp the knowledge presented by the system. *Cognitive or epistemological adequacy* of the representation might be the right term for denoting this desideratum. This also leads us to the fundamental question of knowledge representation of whether there are *primitive representational structures*, richer than just objects, relationships, and so forth, that pervade human knowledge representation.

With anything related to computers, we have to be concerned about *efficiency* in the processes to be carried out by the machine. Otherwise we could spend most of our time waiting for responses from the computer. This again puts a constraint on the representation since an operation is performed at a different speed in one representation than in another.

These last two desiderata are actually constraints on representations in two different media, assuming current computer technology. While cognitive adequacy addresses the representation on the level of the human/machine interface, efficiency plays a role on the level of the machine (i.e., hardware) only. This demonstrates the importance of the operations for our considerations which are very different ones (human interaction vs. operations on hardware) for these two levels. In fact, this also demonstrates the usefulness of talking of different *levels* of representation and of the representational medium. So far we have mentioned two such levels which could be subdivided into further ones according to the kind of user or level of implementation.

3 FORMS OF REPRESENTATIONS

In this section we briefly describe ten of the major forms of representation in use.

3.1 Natural Language Representation

Natural language provides a form of representation that has become the standard medium for communication among people. Sometimes natural language is complemented by drawings, figures, gestures, and the like. None of these qualify under the notion of a scheme as noted in Section 2.1. Nevertheless it is helpful to take into account these forms of representation as a guide for the development of more formal ways of representation.