



Five copies of a draft paper (up to 20 pages) should be sent to John Boose before May 16, 1990. Acceptance notices will be mailed by July 18. Full papers (20 pages) should be returned to the chairman by September 17, 1990 so that they may be bound together for distribution at the workshop.

There will be a travel-and-expense award of up to \$500.00 US for the best paper submitted by a graduate student. Please note if the paper should be considered for this award.

Workshop Co-chairmen:

Send papers via US mail to:

John Boose	Brian Gaines
Advanced Technology Center	Dept. of Computer Science
Boeing Comp. Services, 7L-64	University of Calgary
P.O. Box 24346	2500 University Dr. NW
Seattle, WA, USA 98124	Calgary, Alberta
(206) 865-3253	Canada T2N 1N4
	(403) 220-5901

Send papers via express mail to:

John Boose
Advanced Technology Center, 7L-64
Boeing Computer Services, Bldg. 33.07
2760 160th Ave. SE
Bellevue, Washington, USA 98008

Program Committee:

Jeffrey Bradshaw	B. Chandrasekaran
William Clancey	Michael J. Freiling
Thomas Gruber	Catherine Kitto
Sandra Marcus	John McDermott
Ryszard Michalski	Mark Musen
Bruce Porter	Mildred Shaw

the collection of reprints edited by Brachman & Levesque (1985). Alternatives include "special issues" like the IEEE Computer issue on "Knowledge Representation" (edited by McCalla & Cercone, 1983) which grew into a more advanced (and esoteric) collection edited by Cercone & McCalla (1987). Other notable volumes include Hobbs & Moore (1985) which is an excellent collection of more formal theories of knowledge representation (again a rather more advanced text).

The fact that all the contributions are original is something to set it apart from other texts. It also benefits from the use of authors "in-house" (so to speak) as there is some cross-referencing and a sense of continuity between chapters. Although there is not enough continuity to be fully satisfactory this is a product of the disparate nature of the enterprise rather than a lack effort. Set against the background of existing literature there is clearly a need for an introductory level text surveying the area of knowledge representation (KR).

There are 11 chapters (including the introduction) covering logic, nets, frames, rules, psychological studies, conceptual graphs, control/meta-knowledge, time, functional approaches, expressive power and computability. The chapters vary in length the longest at 44 pages (psychological studies) and the shortest just over 10 pages (expressive power and computability). None of the chapters are bad and some are really quite good. I shall talk mainly about the good ones.

Reasonable descriptions of both logic-based approaches (ch. 2 by Pavelin) and semantic networks (ch. 3 by Mac Randal) are available readily elsewhere - computer scientists usual study the former and psychologists the latter. Here's a chance to get accessible introductions in one volume. There is a chapter on conceptual graphs (a rather more formal semantic net approach) which provides a brief overview of the topic (chapter 7 Jackman & Pavelin).

Chapter 4 on Structured Object Representation - Schemata and Frames (Ringland) tells us how it is from the AI end of the spectrum rather than the psychological end (e.g. frames/schema theory Minsky-style). It is an insightful contribution and made points about the usefulness of such a representation that I now feel much clearer about. Ringland discusses Hayes' attack on the vagueness of Minsky's frame concept and how they could be expressed as first order logic. It was probably this vagueness that led to the frame concept being taken up and applied so freely, but therein lies its strength and weakness. You have a theory that is vague enough that it can seem very powerful - the more precise you make the theory the less powerful it apparently becomes (my comment not Ringland's). Ringland tempers Hayes' criticism without dismissing it outright and produces a nice critical overview of frame-based representations.

Psychological studies of knowledge representation (6, by Conway & Wilson) is a real bumper chapter. It covers a fair size chunk of material from production systems, semantic nets, frames and schemata through to mental imagery, mental maps and mental models. A whole course on all the interesting topics of cognitive psychology in 44 pages! Great stuff. In some ways it acts a unifying chapter for a lot of the material covered in the book. My extra suggestions for further reading include the recently published "Readings in Cognitive Science" edited by Collins & Smith (1988) which has a section on representation and I would also mention Gentner & Stevens' collection - "Mental Models" (1983).

This is not a book on how to use the approaches covered and there are no examples of code (apart from short bits of pseudo code). If you want that kind of "KR by example" ap-

Book Reviews

Approaches to Knowledge Representation - An Introduction

Edited by G. A. Ringland & D. A. Duce
Research Studies Press Ltd.
John Wiley & Sons Inc.
\$65 (Cloth), Approx \$35 (Student Edition)

Reviewed by: Simon Ross
Department of Computer Science
University College London
London WC1E 6BT

JANET: sross@uk.ac.ucl.cs

ARPA: sross@cs.ucl.ac.uk

UUCP: ...ukc!ucl-cs!sross

UUCP: mcvax!ukc!ucl-cs!sross

BITNET/EARN: sross%uk.ac.ucl.cs@ukacrl.bitnet

This collection is entirely edited and authored by members of the Knowledge Based Systems Group, Informatics Division of the Rutherford Appleton Laboratory, Didcot, UK, a research Institute funded by the Science & Engineering Research Council. As such it is an entirely British collection which provides an introductory text with a slightly different viewpoint as nearly all contributions to the field have been of North American origin (at least this is the impression one gets from other texts).

The main source-book previously available in this area is

proach then you could try Winston (1985) for Frames in Lisp, Frost (1986) does a reasonable job on how to get from data to knowledge representation, Charniak et al (1987) is a good source for various approaches in CommonLisp, the trusty collection by Schank & Riesbeck (1981) is good for schema-type approaches, and there are numerous textbooks with examples of prolog (e.g. Bratko, 1987). The most impressive and sustained arguments for production systems as a knowledge representation formalism is found in Klahr et al (1987), the work by Newell and colleagues on the SOAR architecture, and of course Anderson's ACT*.

Conway & Wilson (ch. 6) provide a section on procedural knowledge and discuss the limitations of Anderson's use of production systems as a general architecture of cognition. I would like to have had more on this. They do not mention SOAR or Klahr et al but it is mentioned as further reading in the "Rule-Based Systems" chapter. This chapter (by Williams & Bainbridge) provides an overview that points out some of the advantages and disadvantages of production systems, but otherwise there is nothing new.

I was a little disappointed with the chapter on Representing Time (chap. 9 by Kwong). It does a reasonable job of identifying how we might tackle the topic (what it is, why it is important) and presents some methods to this end. It might have been better mentioning some applications - such as controlling process plants. It doesn't supply any further reading but cites the important work by J. F. Allen but neglects to mention Shoham (1987) which is a very influential recent text in this area.

One area the collection does not cover is the use of Hybrid KR as a pragmatic and realistic solution. Given that not all knowledge fits nicely into, for example, the production system formalism for expert system development - we often need many formalisms to tackle problems of non-trivial complexity. After all why are there so many approaches to KR around? Because any one formalism only fits a (usually small) number of problems. What about a chapter on what types of formalism best fits certain types of problems? A major problem at present is that we have no way of testing the efficiency and efficacy of various KR because no one problem can be adequately represented by enough approaches to assess the computational pay-offs of using one KR over another. What is needed is some form of representation that could allow for these apparently disparate approaches to be pulled together (some people advocate a foundation in logic for this purpose).

Having said that let us consider the last two chapters on the book- these are the type of issues that are addressed, in a rather more specific way. I see these kinds of issues to be fundamental to KR research. Chapter 10, Functional Approaches to KR (Lambert) looks at semantic net/frame type formalisms and matches them to software engineering principles, for example, handling representations, in this case frames, like data structures. The theme is continued by Williams and Lambert who introduce the ideas of expressiveness and computability (after Levesque & Brachman, 1985, in Brachman & Levesque, 1985). The thesis behind this chapter is that "a knowledge representation system must do more than just represent.. it must be modifiable and there is a payoff between the expressiveness of the representation, the complexity of the task and the amenability "to encoding and manipulation within a computer program" (p.223).

Finally, we come to PDP approaches - what no PDP! Well if they don't mention nor will I. (Well it is mentioned in passing in chapter 6).

Although books of this type are undoubtedly aimed at the computing/AI end of the market. I would recommend that cog-

nitive scientists (cognitive psychologists etc) also have a look as it strikes a good balance between the two fields and is digestible for both.

The reasonable price (for the student edition) makes this book an attractive offering - I think it has filled a gap in the market.

Footnote: Since many of the authors are ardent fans of R. J. Brachman and no doubt other readers will be too - check out a recent offering not cited "The Basics of Knowledge Representation" in the AT & T Technical Journal, 67,1, January-February, 1988.

References:

- Brachman, R. J. & Levesque, H. J (1985) Readings in Knowledge Representation. Morgan Kaufmann.
- Bratko, I. (1987) Prolog Programming for Artificial Intelligence. Addison Wesley.
- Cercone, N. & McCalla, G. (1987) The Knowledge Frontier. Essays in the Representation of Knowledge. Springer Verlag.
- Charniak, E., Riesbeck, C. K., McDermott, D. V., & Meehan, J. R. (1987) Artificial Intelligence Programming. LEA. (Second Edition).
- Frost, R. A. (1986) An Introduction to Knowledge Based Systems. Collins.
- Gentner, D. & Stevens, A. L. (1983) Mental Models. LEA
- Hobbs, J. R. & Moore, R. C. (1985) Formal Theories of the Commonsense World. Ablex.
- IEEE Computer (1983) Special Issue on Knowledge Representation (edited by G. McCalla & N. Cercone), 16, 10.
- Klahr, D., Langley, P. & Neches, R. (1987) Production Systems Model of Learning and Development. MIT Press.
- Shoham, Y. (1987) Reasoning about Change. Time and Causation from the Standpoint of Artificial Intelligence. MIT Press.
- Stevens, A. L. & Smith, E. E. (1988) Readings in Cognitive Science. Morgan Kaufmann.
- Winston, P. H. (1985) Artificial Intelligence. Addison Wesley. (Second Edition).

Annual Review of Computer Science

J. F. Traub, ed.

Annual Reviews Inc.

4139 El Camino Way, P.O. Box 10139

Palo Alto, CA 94303-0897

Reviewed by

Faisal Saeed

Dept of Computing and Information Science

Oklahoma State University

Stillwater, OK 74075

saeed@a.cs.okstate.edu

Preparing a review of a field as vast as Computer Science on annual basis is not an easy job. A common problem with such reviews is that the material presented in them lacks the cohesion associated with a book and look more like just a binding of articles. In this regard the Editorial Board of the Annual Review has done a great job by not only selecting papers on active research topics but also presenting them in a book-like manner to provide the reader with an ease of reading.

Two volumes have already been published in 1986 and 1987. The volume in hand, published in 1988, is the third in the series and presents a review of research in different areas in computer science up to 1986. Topics included in this volume are:

- Database Security