American Journal of Computational Linguistics *Microfiche 55 : 14*

REPRESENTATION AND UNDERSTANDING

STUDIES IN COGNITIVE SCIENCE

EDITED BY DANIEL G. BOBROW AND ALLAN COLLINS
Xerox Palo Alto Research Center and Bolt Beranek and Newman

Academic Press, Inc.
New York

LC 75-21630 $15.00 ISBN 0-12-108550-3

REVIEWED BY JOHN MYLOPOULOS

Department of Computer Science
University of Toronto
Toronto, Ontario, Canada M5S 1A7

A major goal of Artificial Intelligence research today is to design systems that "understand" a body of knowledge, i.e. use it whenever appropriate. The representation of the knowledge available to such an "understander" system is an important issue for the system's design and is intimately related to the proposed uses of that knowledge. This book includes a collection of thirteen papers written by some of the best known researchers who are currently working on understander systems. The papers were selected among those presented at a conference held in memory of Jaime Carbonell.

The contents of the book are as follows

I. Theory of Representation

1. Dimensions of Representation

   Daniel G. Bobrow

2. What's in a Link: Foundations for Semantic Networks

   William A. Woods

3. Reflections on the Formal Description of Behavior

   Joseph D. Becker

4. Systematic Understanding:
   Synthesis, Analysis, and Contingent Knowledge
   in Specialized Understanding Systems

   Robert J. Bobrow & John Seely Brown

II. New Memory Models

5. Some Principles of Memory Schemata

   Daniel G. Bobrow & Donald A. Norman

6. A Frame for Frames:
   Representing Knowledge for Recognition

   Benjamin J. Kuipers

7. Frame Representations and
   The Declarative-Procedural Controversy

   Terry Winograd

III. Higher Level Structures

8. Notes on a Schema for Stories

   David E. Rumelhart

9. The Structure of Episodes in Memory

Roger C. Schank

10. Concepts for Representing Mundane Reality in Plans

Robert P. Abelson

IV. Semantic Knowledge in Understander Systems

11. Multiple Representations of Knowledge for Tutorial Reasoning

John Seely Brown & Richard R. Burton

12. The Role of Semantics in Automatic Speech Understanding

Bonnie Nash-Webber

13. Reasoning From Incomplete Knowledge

Allan Collins, Eleanor H. Warnock, Nelleke Aiello, & Mark L. Miller

As stated in the book's introduction, the section on "Theory of Representation" deals with general issues regarding the representation of knowledge, while that on "New Memory Models" discusses the implications of the assumption that input information is always interpreted in terms of large structural units derived from experience. The section titled "Higher Level Structures" focuses on the representation of plans, episodes and stories within memory. Finally, the section on "Semantic Knowledge in Understander Systems" describes on-going work of the SOPHIE, SPEECHLIS and SCHOLAR projects at BBN.

In attempting to review the papers that appear in this book collectively rather than individually, we arrived at a slightly different taxonomy than that used by the book's editors. For the discussion that follows, numbers from 1 to 13 refer to the papers in the book. Other references, numbered from 14 on, are given at the end of the review.

1. Comparison of Representations and General Criteria for their Adequacy.

Many different paradigms have been proposed for representations, including ones based on Predicate Calculus, Production Systems, Semantic Networks, frames", and ones that are PLANNER- or ACTOR-like. Winograd [14] provides an excellent comparison of those paradigms. An important and controversial aspect of current work on representation is the debate on the distinction between and desirability of declarative vs. procedural representations, and episodic vs. semantic memory organizations

A number of papers either provide criteria for comparison among different representations or general frameworks within which these representations can be described and discussed. Others present adequacy criteria for representations or discuss one or more of the controversies mentioned in the previous paragraph.

[1] proposes several "dimensions" along which representations can be compared and illustrates the use of those dimensions for the comparison of three very simple representations for digitized binary pictures. This paper also serves as an extended introduction to the rest of the book.

Becker describes in [3] how computer science concepts such as scheduling, backtracking, interrupts etc. can be used to model aspects of (human or machine) behaviour such as goals, conflicts, spheres of influence and decision making. Although the paper does make several interesting points, the lack of rigor hurts the discussion. For example, the last section of the paper presents an argument in favour of the view that behavioural descriptions are relative in the sense that behaviour admits many different, and possibly ambiguous, descriptions, unlike, say, a capacitor charging or discharging. But surely one could argue that the capacitor's behaviour could also admit different and ambiguous descriptions, such as "the capacitor is delaying a signal", "the capacitor is filtering out certain undesirable frequencies" etc. If one accepts this view, then there is no straightforward, absolute, canonical or true description for anything, not just for behavioural systems. Perhaps the author is trying to establish a different point. If so, we missed it.

[4] presents the SCA model, which is intended to provide a framework for designing and comparing understander systems. The discussion gives accounts of two modules that are part of the model, the first to integrate incoming information to the system's knowledge base, the second to use the knowledge base in order to answer questions. Three existing systems are described within the framework of the SCA model as evidence of the model's adequacy. As admitted by the authors, however, the model is a very partial answer to an overall organization for

a system involving many processes. It should be added that with the discussion being so general and devoid of detail, it is hard to see whether a genuine contribution is being made or whether the model's apparent ability to fit different existing systems is precisely due to the lack of detail.

The paper by D Bobrow and Norman [5], proposes (memory) schemata as the constructs in terms of which the organization and operation of a memory can be described. The properties schemata should have are then discussed and many requirements are set forth for the adequacy of a representation Some of these are the use of context-dependent descriptions to access schemata, the accountability of all inputs, i.e. the ability of a memory system to account for all inputs, no matter how trivial, at some level, and the distinction between data-limited and resource limited processes. The overall framework that emerges is quite interesting because it takes into account issues regarding the design of large resource-limited systems that had only been studied in the past in Operating Systems literature.

The first part of Winograd's paper [7] deals with the declarative vs. procedural representation controversy and the trade-offs involved. The controversy is an old one within computer science and includes, among other things, the merits and demerits of a (declarative) representation that allows programs to be represented as data. The discussion in the paper is quite well-written and argues convincingly that the basic trade-off between the two different types of representations is one of modularity, for declarative, vs. flexible interaction among different facts, for procedural.

Schank's paper [10] includes a discussion on whether the organization of human memory is episodic or semantic. An episodic memory organization implies that knowledge is stored as temporally dated episodes and events, with temporal spatial relations linking these events. A semantic memory organization, on the other hand, involves time-invariant knowledge a person possesses, e.g., "all elephants are animals". A corollary of these definitions is that an episodic memory organization favours temporal and causal connectives (e.g., THEN, REASON, ENABLE etc.), whereas a semantic memory organization uses extensively the "ISA hierarchy" (e.g., "an elephant is-a animal"). The discussion presented in the paper on this issue is somewhat confusing since at one point (pp. 255-256) the two types of organization are contrasted as if they were mutually exclusive, while later on (p. 263) the paper argues for a combination of the notions of semantic and episodic memory. In either case, Schank's work certainly makes a convincing argument in favor of an episodic memory organization by showing how it can be used to represent the meaning of a paragraph.

II. Critique and Extensions of Representation of Knowledge Paradigms

Several papers, including some that were mentioned in the previous section, criticize, refine, or extend one of the existing paradigms for the representation of knowledge.

The most notable example among those in this category is Woods' paper [2] which criticizes many (mis)uses of semantic networks by pointing out situations where their semantics are poorly defined or even inconsistent. Particular attention is paid to the representation of quantification and that of relative clauses.

As many of the readers undoubtedly know, Minsky's influential paper introducing "frames" [15] provides more of an ideology than a theory for representing knowledge. Kuipers in [6] argues in favor of a number of properties frames should have, such as the ability to describe an object or situation to varying degrees of detail, the ability to be instantiated and the ability to handle small perturbations of expected input data without major failures. He illustrates the desirability of these features with a simple example of object recognition.

The second half of Winograd's paper makes an attempt to synthesize declarative and procedural aspects of a representation. His proposal is based on frames and uses a generalization (ISA) hierarchy having a number of features, including the ability to associate procedures to objects on the hierarchy which specify how to perform different operations on those objects. Many of the ideas in [5] and [7] have been incorporated in KRL [16], as developed by D. Bobrow and Winograd.

III . Representing Different Kinds of Knowledge

Information entering an understander system may have many different "forms", i.e. it may be coded as photographs or line drawings, simple sentences or paragraphs or even complete stories. Moreover, it may have different "content" i.e. involve a fairy tale world of kings and dragons, a blocks world of cubes and pyramids, a social, mental or physical world. One important aspect of the representation problem is the definition of a collection of knowledge, defined by a restriction on its form and/or content, and the investigation of the adequacy of a particular representation.

As mentioned earlier, Woods' paper does discuss the representation of quantification in terms of semantic networks, where the form of the knowledge involved is presumably (first order) Predicate Calculus and the content is unconstrained. It also discusses the representation of relative clauses and complex sentences where the form is natural language and the content is, again, unconstrained.

Rumelhart's paper [8] is primarily concerned with the discovery of structure underlying simple stories. The structure is defined in terms of a phrase structure grammar with semantic rules associated to each production. The paper certainly follows the general trend towards studying linguistic units larger than sentences, such as paragraphs, dialogues or stories. Whether the methodology used (in particular, phrase structure grammars) will be found adequate for the description of structure in stories remains to be seen.

Schank [9] deals mainly with the problem of constructing a structure of causally-linked actions and changes of states (episodes) from a paragraph. When episodes are used to make sense of new inputs in often-experienced situations, they are called "scripts". The paper ends with a brief introduction of scripts. More details about them can be found in more recent publications by Schank and his students, e.g. [17,18].

Rumelhart's and Schank's work are related in that they both attempt to define the structure of a collection of knowledge limited with respect to form (stories for Rumelhart, paragraphs for Schank) and unconstrained with respect to content. Moreover, both papers agree that the underlying representation used must involve causally-linked events, and the causal connectives they employ are similar.

Abelson's paper is concerned with the representation of "mundane reality" involving social actions. The approach he follows is to postulate a number of primitive states and actions for achieving these states, in terms of which hopefully all simple social behaviour can be described. The discussion of the primitives is quite thorough, but the examples given do not provide sufficient evidence that the primitives proposed are in fact descriptively adequate. Abelson's work is complementary to Schank's in several respects and there is more recent joint work on the subject [19].

IV'. On-going Projects involving Understander Systems

The last three papers of the book discuss particular projects involving the design and implementation of understander systems.

[11] describes the scope, basic methodology, and achievements of SOPHIE, a knowledge-based computer aided instruction (CAI) system which attempts to teach a student about electronic circuits by asking questions, answering questions and letting him try out his ideas. Of particular interest to computational linguists should be the section describing the "semantic grammar" developed by Burton to handle the types of sentences expected during a dialogue on electronic circuits.

Nash-Webber [12] provides an overview of the BBN SPEECHLIS project in the context of a discussion on the use of semantic knowledge for speech understanding. Finally, [13] discusses some of the inference rules implemented or being considered for implementation by the SCHOLAR project whose aim is to develop a knowledge-based CAI system that teaches geography. The reader may find many of the rules stated in the paper completely reasonable and yet quite shaky from a logical point of view. For example, one rule (the uniqueness assumption) states that if only one thing is found, it can be assumed that it constitutes a complete set. Thus if someone knows of only one city called "Springfield" and located in Massachussetts, he can use the uniqueness assumption to reply "no" to "Is Springfield in Kentucky?" even though there may well be such a city.

The papers in this section constitute an important complement to the rest of the book which often involves discussions that are too far removed from the reality of an implemented (or implementable) system.

Overall, this book provides an excellent review of the state of the art, circa 1975, on the problem of representing knowledge.

•It should be apparent from the previous discussion that the book assumes a familiarity with basic issues of representation and understander system design. For more introductory discussions, the reader is referred to [14] or Schank and Colby [20].

References

14. Winograd, T. "Five Lectures on Artificial Intelligence" Stanford AI-Memo 246, September 1974.

15. Minsky, M. "A Framework for Representing Knowledge" in Winston P. (Ed.) The Psychology of Computer Vision, McGraw Hill, 1975.

16. Bobrow, D. and Winograd, T. "A KRL User's Manual" (unpublished).

17. Schank, R. "Using Knowledge to Understand" TINLAP Proceedings pp. 117-121, June 1975.

18. Schank, R. and the Yale AI Group "SAM --- a Story Understander" Yale University, Dept. of Computer Science, August 1975.

19. Schank, R. and Abelson, R. "Scripts, Plans and Knowledge" Proceedings IJCAI, pp. 151-157, September 1975.

20. Schank, R. and Colby, K. (Eds.) Computer Models of Thought and Language, Freeman, 1973.