

system as more than a mere existence proof: he offers it as a psycholinguistic theory of how children actually acquire their first language, and he frequently cites observations about child language as tending to confirm his theory.

A book that can make even a *prima facie* plausible claim to have achieved these things must be an important one; hence the length of this review. Just how far Berwick's attempt has succeeded is a question not to be answered quickly. The book is dense, and Berwick is not always as skilled as he might be at helping the reader to disentangle the central skeleton of his exposition from peripheral technical details. An adequate assessment of Berwick's work will need extended consideration by the scholarly community, preferably with further elucidation by Berwick himself.

To set the ball rolling, let me mention some points that worried me on a first reading of the book, though I do so without any assumption that they will ultimately prove fatal to Berwick's case.

My chief worry is that Berwick is quite vague about the computational implementation of his system, and about the degree of success it has attained in practice. The target he sets for his system is to reconstruct the set of rules contained in Marcus's parser, which Berwick describes as numbering "on the order of 100". At one point Berwick states that "by the time the system has processed several hundred sentences, it has acquired approximately 70% of the parsing rules originally hand-written for [Marcus's] parser". Later, he says that "On the order of 70-100 rules . . . are learned from a corpus of several hundred sentences". What is the truth: does the system achieve 70%, or 100%, success with respect to the chosen criterion? If the former, what sort of errors are made? How far does the number of data sentences required tend to vary with the order of data presentation? How naturalistic are the data?

Secondly, although Berwick frequently claims that his theory makes correct predictions about child language, he is again vague about the facts in this domain: ". . . the NP after *by* could be taken, incorrectly, as the Object of the verb. This seems to happen with children"; ". . . children seem to frequently drop Subjects . . . Hyams (1983) has confirmed this . . ." (referring to an unpublished paper which is not discussed further; such remarks are characteristic, and unsatisfying).

This latter point is the more worrying, since it is clear in other respects that Berwick's expertise lies more in the computational field than in the natural language domain. Occasionally he perpetrates linguistic howlers; for instance, he claims that the ungrammaticality of **There were a riot on Tuesday* shows that existential *there* governs subject-verb agreement, and he describes the word *assigned* as "trisyllabic". Berwick's weakness in the area of empirical linguistics has the consequence that he is excessively willing to accept every temporary theoretical proposal of the M.I.T. school of linguists as gospel, fail-

ing to distinguish long-term, core principles from trial balloons which someone floated last year and which will probably be abandoned next year. In Part II of his book, which discusses the theoretical implications of the acquisition model as opposed to its internal workings, Berwick commits himself to a number of linguistic beliefs that seem indefensible. He spends some time discussing a constraint on natural-language semantics (attributed to F. Sommers and F. Keil) according to which graphs representing the relationship of predictability between vocabulary items rarely or never contain *M*-shaped subgraphs. This seems quite wrong (a *kitten* and a *baby* may both be *stillborn*, a *baby* and an *engineer* may both be *British citizens*), and it is not clarified by the diagram Berwick uses to illustrate it (p. 270). He devotes many pages to arguing that his theory explains an alleged phonological constraint, quoted from an unpublished doctoral dissertation by M. Kean, which (as Berwick describes it) appears to forbid the occurrence of a language having stop consonants at three or more places of articulation but a fricative at only one. In reality, systems with /p t k/ and /s/ but no /f/, /x/, etc. are rather common.

These aspects of Berwick's book seem regrettable, and unnecessary. The book would have been a significant contribution if most of Part II had been omitted. I wish it had been, and that the computer implementation had been discussed more fully; but I hope these points will not cause the valuable parts of the book to be neglected.

Geoffrey Sampson

Department of Linguistics and Phonetics
University of Leeds
Leeds LS2 9JT England

References

- Fodor, J.A.; Bever, T.G.; and Garrett, M.F. 1974 *The Psychology of Language*. McGraw-Hill.
 Marcus, Mitchell P. 1980 *A Theory of Syntactic Recognition for Natural Language*. MIT Press.
 Wexler, Kenneth and Culicover, Peter W. 1980 *Formal Principles of Language Acquisition*. MIT Press.

STRUCTURES AND PROCEDURES OF IMPLICIT KNOWLEDGE (Advances in discourse processes 17)

Arthur C. Graesser and Leslie F. Clark,

Norwood, NJ: Ablex Publishing Corp, September 1985,
viii+326 pp.
ISBN 0-89391-192-5, \$42.50 (cloth); ISBN
0-89391-362-6, \$24.50 (pbk)

How does one understand a narrative? The current scientific theories reduce comprehension to the generation of a correct sequence of inferences from some knowledge structures. Inferences are obviously created when individuals comprehend text, but there is widespread disagreement about what inferences are generated, when

inferences are generated, how many inferences are generated, and what knowledge sources contribute to the generation of inferences.

In their book *Structures and Procedures of Implicit Knowledge*, Graesser and Clark, two psychologists, attempt to answer these questions by presenting a model of comprehension that primarily focuses on knowledge-based inferences (viz. products of what the comprehender knows about the world).

The likelihood of a particular inference depends on the content of the inference, together with (a) the text's context, (b) world knowledge structures and inference engines that are available, (c) the goals of the comprehender, and (d) the pragmatic context of the communication act. Admittedly, this is too much to handle at once, and the authors do not wait until the last section of the book to clearly state the goals and limitations of their work. There is no discussion of syntactic parsing, no formal theory of meaning, and no pragmatic model. This is not a book about linguistics but rather about conceptual modeling and, specifically, about the generation and usage of knowledge-based inferences during text comprehension. The authors are quick to point out that formal work in linguistics, logic, and philosophy, as well as AI research, ignore "important characteristics of human cognition". Implicitly, researchers in those fields are invited to momentarily leave their idealistic vacuums or their Lisp code in order to refresh their knowledge about text comprehension and psychological plausibility.

Though the book does not present an exhaustive survey of the available research on inferences and text comprehension, its concise discussion in the first chapter of inference taxonomies and engines and its rich bibliography make it an excellent reference.

Text comprehension is extremely complex, and the authors can only offer a very partial yet quite interesting solution: they propose procedures to model comprehension, recall, summarization, and question answering. These procedures work on generic knowledge structures (represented by conceptual graphs) that they traverse and match in order to generate the inferences that make the text coherent, as well as other inferences that capture the comprehender's expectations. From a computational point of view, since there is no implementation of the model, the discussion may sometimes appear superficial. Also, Graesser and Clark too often claim without any further explanation that their model includes previous work. As is usually the case for this domain of research, since the stories analyzed must minimize the role of the components left out of the model, the reader is confronted with truly artificial and insipid texts.

Finally, be forewarned! The authors present a generous amount of statistics obtained from numerous experiments; a great deal of time is spent analyzing the data and defending the methodology. The reader may often want to skip to the end of chapters, where good summaries of the results and conclusions are provided.

Jean-Pierre Corriveau
Department of Computer Science
University of Toronto
Toronto, Ontario
Canada M5S 1A4

**CONCEPTUAL STRUCTURES: INFORMATION PROCESSING
IN MIND AND MACHINE**
(The systems programming series)

John F. Sowa

Reading, MA: Addison-Wesley, 1984, xiv+481 pp.
ISBN 0-201-14472-7

John Sowa has written an excellent book. It is beautifully written, and presents a clean, precise look at knowledge representation and its applications. The book combines a sweeping historical perspective from the ancients to current research, with a formal definition of knowledge representation structures.

The first two chapters provide the motivation and set the tone for the rest of the book. They are a delight to read. Chapter 1, "Philosophical Basis", shows why psychologists, linguists, philosophers, and computer scientists are all interested in the problem of knowledge representation, and how their perspectives on the problem overlap and differ. One finishes the chapter with an understanding of the historical development in each of the areas, and the interdisciplinary nature of the field. Throughout, Sowa stresses the importance of formal models, as opposed to ad hoc solutions. Chapter 2, "Psychological Evidence", surveys numerous psychological experiments that illustrate the nature of human language behavior. Modeling this behavior is the essential problem of AI research. The chapter contains several wonderful anecdotes – for example, the eidetic-memored Shereshevskii who lost his job as a newspaper reporter because he could not abstract from detail.

At Chapter 3, "Conceptual Graphs", we enter the technical part of the book. This chapter is a fine introduction to semantic net representation, which he calls conceptual graphs. Conceptual graphs are a canonical form of many AI knowledge representation schemes. In the form he follows throughout the book, Sowa first presents a general, understandable discussion of what is to be represented and why, and follows with formal definitions. He covers all the fundamentals: concepts, generalization and specialization, types and tokens, aggregation and individualization. By the end of the chapter we have a good intuitive understanding of the representation, and a sound formal basis.

Chapter 4, "Reasoning and Computation", proceeds naturally from Chapter 3. We see larger organizations of memory structures, and how they are used in reasoning processes. He shows how logic can be represented in graphs and how deduction can be performed on them.