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-memoried 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.

Chapter 5 is on "Language". Again, we have an excellent, broad summary: child language abilities, strata of language, case grammar, and generation. There is a nice presentation of syntactic analysis. However, I would like to see a deeper discussion of ATNs; here he just mentions the phrase. The sections on context, and integrating syntax and semantics are particularly well done.

Chapter 6, "Knowledge Engineering", is a bit of a grab bag. On expert systems, he mercifully avoids the current hype by taking a "Just the facts, Ma'am" attitude. The other sections on natural language systems, database semantics and inference, knowledge acquisition, and learning are all informative, succinct descriptions of basic problems and the current state of the art.

Chapter 7, "Limits of Conceptualization", ends the book with a description of the wonderful aspects of human behavior that computers can't do at all. It is great fun to read.

The potential audience for this book is extensive. Because the chapters are organized with introductory material at the beginning of each section and formal descriptions later, readers may skip sections without losing the thread of the book. A newcomer to AI can read the first parts of each section, serious students can read it all.

As a textbook, *Conceptual Structures* has both positive and negative aspects. On the positive side, it gives an historical perspective of many fields, presents a general, formal definition of knowledge representation and its uses, and contains good exercises. An additional bonus is the conceptual catalog given in Appendix B. This is exactly what students always ask for: a collection of structures that can be used by programs. On the negative side, specific AI systems are nowhere discussed in detail. It would be necessary for the instructor to provide supplementary material and original sources.

I agree with Clancy (1985) that "Every AI and Cognitive Science researcher should study the conceptual graph notation and understand its foundation in logic, database, and knowledge representation research."

Sharon Salveter
Department of Computer Science
Boston University
111 Cummington Street
Boston, MA 02215

Reference

Clancy, W.J. 1985 Review of *Conceptual Structures: Information Processing in Mind and Machine* by J.F. Sowa. *Artificial Intelligence* 27(1): 113-124.