On *Conceptual Structures*, a response to the review by S.W. Smoliar

Response by: John F. Sowa
IBM System Research, 500 Columbus Avenue, Thornwood, NY 10594, U.S.A.

In his review [28], Stephen Smoliar praised my 1976 article on database query, but he missed the point of my book. He thought that it was constrained by "Codd's relational model of data bases" instead of Brachman's "broader questions of semantics". Yet the database material is part of my attempt to build bridges from AI to other disciplines; I presented even more material on philosophy, linguistics, and cognitive psychology. In continuing his criticism, Smoliar mentioned Brachman's article [2] "which would lead one to question just how far one could go with the elegant simplicity of Sowa's foundation." In fact, the "elegant simplicity" is largely derived from the existential graphs of the philosopher Charles Sanders Peirce. Peirce was too good a logician to commit any of the fallacies that Brachman described.

In the book, I set out to develop a framework for semantics based on a synthesis of AI and cognitive science. On page 1, it asks the question of what is knowledge. On page 2, it gives a preliminary answer: "Knowledge is more than a static encoding of facts; it also includes the ability to use those facts in interacting with the world. A basic premise of AI is that knowledge of something is the ability to form a mental model that accurately represents the thing as well as the actions that can be performed by it and on it. Then by testing actions on the model, a person (or robot) can predict what is likely to happen in the real world." Page 4 states the basic hypothesis: "This book develops the theory of conceptual graphs as a method of representing mental models, shows how it explains results from several fields, and applies it to the design of more intelligent, more usable computer systems." Chapter 1 analyzes the philosophical implications of that hypothesis, Chapter 2 presents psychological evidence for it, Chapter 3 develops the formalism of conceptual graphs as data structures for representing knowledge, Chapter 4 presents their rules of inference and model theoretic basis, Chapter 5 shows how they relate to language, Chapter 6 applies them to knowledge based systems, and Chapter 7 considers the limitations of conceptual thinking and symbolic systems in general. Somehow, Smoliar missed these points. He complained that "Sowa shows a general tendency to get his syntactic act together and take it on the road before considering whether or not he should bring along some semantics as well."

Ideally, technical issues should be settled by reasoning. But since Smoliar's main points are all matters of opinion, the best way to answer them is to quote other reviewers who had different opinions:

- William Clancey [3]: "No other AI text achieves so much in breadth, style, and mathematical precision. This is a book that everyone in AI and cognitive science should know about, and that experienced researchers will profit from studying in some detail."
- Chris Riesbeck [24]: "The breadth of material covered, and (in the central chapters) the depth to which the material is treated, is amazing. There are many surveys of AI in print, and a handful of volumes on programming techniques, but very few theoretically rich textbooks. Sowa has managed to produce one. I recommend it to all students interested in knowledge representation."
- Sharon Salveter [25]: "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."
- John Fox [6]: "This has the feeling of a mature work, a long time in the making and reflecting years of serious thought. It is surprising therefore that the book includes chapters on recent topics such as knowledge engineering as well as classical ones. This is one of its strengths as Sowa is able to cut through some of the AI jargon and relate the new ideas to established topics such as databases and query languages. The references show the same careful mix; they are comprehensive but up to date throughout . . . . It would be on my list of six AI books to take to a desert island, and any student of cognitive science will find it very rewarding."

Each of these reviewers is actively pursuing research in areas related to the book: Clancey in expert systems, Riesbeck in natural processing, Salveter in both natural language and database systems, and Fox in cognitive psychology.

Smoliar also made a number of detailed points that can be answered briefly:

- He complained that the book is "out of date", since it was conceived in the early 1970s. In fact, I discarded or rewrote a great deal of what I had written before 1976; fully half of what appears in print was written from 1981 to 1983. What led me to revise my earlier writing was not some breakthrough in AI itself, but my discovery of Peirce's existential graphs, which had been published in 1896. Timeliness has little correlation with publication date.
- On expert systems, Smoliar says "The contribution here is extremely weak." Yet Clancey, who has been working on expert systems at Stanford for many years, felt that its contribution was significant: "I realized that this book had completely changed my idea of what knowledge representation is. Rather than thinking in terms of 'attributes' and 'values', I started to think in terms of concepts described in relation to other concepts where relations themselves are typed and related to more primitive relations. These ideas have been

in various circles of AI for a decade, but until I read this book, I didn't understand their relevance to heuristic, rule-based programs."

- Smoliar claims that I am unfamiliar with the "knowledge representation hypothesis" and quotes a passage from a dissertation by Brian Smith. Yet Smith's point is a paraphrase of my discussion of knowledge and models on page 2.
- Smoliar claims that "Sowa is trying to address questions of knowledge representation by drawing pictures" and "evade any commitment to a set of questions to be addressed". Yet the first two chapters explicitly state the problems and constraints that a knowledge representation must handle. Then Chapters 3 and 4 develop an abstract mathematical system that meets those constraints. The book states that no aspect of the theory depends on the notation in any way; the graph diagrams are informal illustrations. In fact, it is even possible to use KL-ONE notation or Schankian conceptual dependency diagrams as a concrete notation for the theory up to Section 3.6. Beyond that section, neither KL-ONE nor CD diagrams can represent the contexts and quantifiers.
- In another comment, Smoliar says "Of course, there is much more to knowledge representation than Sowa's view of it." Yet he fails to mention anything specific. I can only respond with a list of what the book contains: a complete system of logic including modal and higher-order forms; a framework that subsumes other AI systems as special cases, including Schank's conceptual dependency theory, KL-ONE, and KRYPTON; procedural attachments for relating the propositional graphs to computational mechanisms; the use of schemata for handling script-like and frame-like reasoning; model theoretic basis including open-world and closed-world models; a theory of how the graphs are related to language and perception; a conceptual catalog that provides a taxonomy of all the concepts and relations used in the examples of the book; a technique for doing conceptual analysis to extend and refine the taxonomy; and in the final chapter, a critique of the problems and limitations of symbolic systems in general and conceptual graphs in particular. That is not all of AI, but it's a lot for 481 pages.

While responding to Smoliar's criticisms, I would like to mention another point he made in a note to me. It was a quote from an unnamed Yale graduate: "Why didn't you come right out and say that there is no difference between Sowa's conceptual graphs and Schank's conceptual dependency?" That question has two presuppositions: that there is no difference and that I have appropriated Schank's ideas without giving proper credit. To address the first point, conceptual graphs go beyond conceptual dependency theory in the following ways:

- They form a complete system of logic, including modal and higher-order forms.

- They include a graph version of the lambda calculus that enables new concept and relation types to be defined and manipulated within the system. (This differs from the early CD theory, which was committed to a fixed set of primitive acts. Those acts could be represented in conceptual graphs, but the definitional mechanisms allow greater flexibility.)
- They include a way of representing contexts that is isomorphic to the contexts in Kamp's discourse representation structures [17]. (This point, by the way, distinguishes conceptual graphs from Hendrix's partitioned nets [14], which do not support Kamp's contexts.)
- Finally, the system is presented in formal assumptions, definitions, theorems, and proofs that explicitly define its capabilities and limitations.

The second point is whether Roger Schank was given proper credit. In May 1968, I developed the first version of conceptual graphs in a term paper for Marvin Minsky's AI course at MIT. At that time, Quillan's semantic memory [20] and Hays' dependency grammar [13] were the only previous graph systems that I was aware of. In their first published paper, Schank and Tesler [26] also cited Hays and adopted the word "dependency" from him. Hays, in turn, derived his graphs from Lucien Tesnière, whose book [34] is a rich source of examples and analyses in graph form. Although my initial ideas were developed independently of Schank's, the book was certainly influenced by him, and the index lists 21 citations to his writings; the only author I cited more frequently is C. S. Peirce.

On the procedural-declarative controversy, I have been strongly influenced by logic programming—see, for example, my chapter on PROLOG in [35]. Conceptual graphs may be viewed as the basis for a logic programming system that is designed to support natural language semantics as directly as possible. Smoliar noted that the graphs are purely propositional and criticized the "imbalance between propositional and behavioral accounts". He correctly observed that a practical system needs more than rules of inference: "the behavioral issue of determining how to perform such inferences remains a problem of imposing magnitude". In Section 4.7, I sketched a method of using the graphs in a conceptual processor. Admittedly, that sketch was vague and incomplete, since I did not have a working system. Since then, however, there have been a number of implementations around the world:

- Fargues et al. [5] implemented a PROLOG-like inference engine using conceptual graphs instead of simple predicates. This provides an important extension beyond PROLOG because the type labels on the concept nodes support a version of multisorted logic and the conclusion of a rule can be an arbitrarily large graph instead of just a single predicate.

- Beringer [1] used a compact encoding of the graphs in a high-performance inference engine. He implemented frame-like inheritance in addition to the

PROLOG-like backward chaining. Since inheritance can also be done with PROLOG rules, implementing it as a rule of inference does not increase the logical power, but it does improve the performance. On Schubert's steamroller problem, Beringer's system executed only 14 inference steps in 8 milliseconds of CPU time on an IBM 3090.

- At Deakin University in Australia, Garner and Tsui [9] implemented an Extendible Graph Processor (EGP) along the lines sketched in Section 4.7. Unlike the logic-based systems of Fargues and Beringer, EGP uses heuristics for frame-like or script-like inferences on conceptual graphs. It has proved to be a versatile tool, which seven graduate students are currently using or extending in their dissertation projects.
- At Sydney University, Rao and Foo [23] implemented CONGRES, a conceptual graph inference system that supports modal reasoning as well as reasoning about knowledge and belief. They are extending it to include DYNABELS, a dynamic belief revision system that supports a version of truth maintenance.
- At the University of Minnesota, James Slagle is leading a group of researchers and graduate students from the Minneapolis-St. Paul area in implementing conceptual graphs in COMMONLISP. Slagle and Gardiner [27] have designed an expert system shell based on conceptual graphs that extends an older system Slagle had implemented with a different version of semantic networks.

These are just a sample of the projects using conceptual graphs. Unlike hybrid systems, which use different languages for different kinds of inference, conceptual graphs can support everything within a single formalism. Furthermore, systems like Beringer's show that they can do so without any loss in performance.

Finally, Smoliar would like to put my book on the list of dangerous things that should be kept away from impressionable young minds: "Greener students who may still be susceptible to gullibility would do well to steer clear". For the truly gullible, a little knowledge is always a dangerous thing. But the only cure for gullibility is wide exposure to as many views as possible. For that reason, I tried to include a wide range of views from ancient philosophy to modern AI. If that range is not sufficient to cure gullibility, Chapter 7 presents a critique of conceptual thinking that Riesbeck found "a surprisingly strong criticism of what has gone before". I didn't want any reader to assume that all the problems of AI and cognitive science have been solved.

REFERENCES

The following references include the papers cited above as well as other papers about conceptual graphs and their applications. The forthcoming collection edited by Sowa, Foo, and Rao presents applications to natural knowledge processing, expert system shells, knowledge acquisition, and database design.

1. Beringer, H., Programming in logic with conceptual graphs, unpublished report, Electronique Serge Dassault, St. Cloud, France. Revised version to appear in [32].
2. Brachman, R., What's in a concept, *Int. J. Man-Mach. Stud.* **9** (1977) 127–152.
3. Clancey, W. J., Review of *Conceptual Structures* by J. F. Sowa, *Artificial Intelligence* **27** (1985) 113–128.
4. Coombs, M. J., Hartley, R. T. and Stell, J. G., Debugging user conceptions of interpretation processes, in: *Proceedings AAAI-86*, Philadelphia, PA (1986).
5. Fargues, J., Landau, M. C., Dugourd, A. and Catach, L., Conceptual graphs for semantics and knowledge processing, *IBM J. Res. Dev.* **30**(1) (1986) 70–79.
6. Fox, J. T., Search for cognitive science, *Nature* **310** (1984) 437–438.
7. Garner, B. J. and Tsui, E., Knowledge representation for an audit office, *Aust. Comput. J.* **17**(3) (1985) 106–112.
8. Garner, B. J. and Tsui, E., Recent advances in computer audit research, *EDP Auditor* **4** (1985) 3–16.
9. Garner, B. J. and Tsui, E., An extendible graph processor for knowledge engineering, in: *SPIE Proceedings, Applications of Artificial Intelligence*, Orlando, FL (1986).
10. Garner, B. J., Lui, D. and Tsui, E., On better goal type interpretation, in: *Proceedings First Australian AI Congress*, Melbourne (1986).
11. Garner, B.J. and Tsui, E., Control scripts for the reorganisation of knowledge bases, in: *Proceedings First Australian AI Congress*, Melbourne (1986).
12. Garner, B. J. and Tsui, E., A self-organizing dictionary for conceptual structures, in: *Proceedings Applications of AI V*, Orlando, FL (1987).
13. Hays, D. G., Dependency theory: A formalism and some observations, *Language* **40**(4) (1964) 511–525.
14. Hendrix, G. G., Encoding knowledge in partitioned networks, in: N. V. Findler (Ed.), *Associative Networks* (Academic Press, New York, 1979) 51–92.
15. Hurwitz, A. and Marshall, J., An intelligent help system for IBM VM/SP, in: *Proceedings SHARE 68*, Session G204, San Francisco (1987).
16. Jackman, J. K., Inference and the conceptual graph knowledge representation language, in: S. Moralee (Ed.), *Research and Development in Expert Systems IV, Proceedings of Expert Systems '87* (Cambridge University Press, Cambridge, 1987).
17. Kamp, H., Events, discourse representations, and temporal references, *Languages* **64** (1981) 39–64.
18. Poesio, M., An organization of lexical knowledge for generation, Tech. Rept., Project WISBER, University of Hamburg, 1987.
19. Poesio, M. and Rullent, C., Modified caseframe parsing for speech understanding systems, in: *Proceedings IJCAI-87*, Milan (1987).
20. Quillan, M. R., Semantic memory, Rept. AD-641671, Clearinghouse for Federal Scientific and Technical Information, 1966. Abridged version in: M. Minsky (Ed.), *Semantic Information Processing* (MIT Press, Cambridge, MA 1966) 227–270.
21. Rao, A. S. and Foo, N. Y., Modal horn graph resolution, in: *Proceedings First Australian AI Congress*, Melbourne (1986).
22. Rao, A. S. and Foo, N. Y., DYNABELS: A dynamic belief revision system, Basser Dept. of Computer Science, Tech. Rept. 301, Univ. of Sydney, Sydney, Australia, 1986.
23. Rao, A. S. and Foo, N. Y., CONGRES: Conceptual graph reasoning system, in: *Proceedings Third IEEE Conference on AI Applications*, Orlando, FL (1987).
24. Riesbeck, C., Review of *Conceptual Structures* by J. F. Sowa, *ACM Comput. Rev.* **27**(8) (1986) 392–393.
25. Salveter, S., Review of *Conceptual Structures* by J.F. Sowa, *Comput. Linguistics* **12** (3) (1986) 218–219.
26. Schank, R. C. and Tesler, L. G., A conceptual parser for natural language, in: *Proceedings IJCAI-69*, (1969) 569–578.

27. Slagle, J. R. and Gardiner, D. A., Athena: An expert system shell using conceptual graphs, Unpublished report, University of Minnesota, 1987. Revised version to appear in [32].
28. Smoliar, S.W., Review of *Conceptual Structures* by J.F. Sowa, *Artificial Intelligence* **33** (1987) 259–266.
29. Sowa, J. F., Conceptual graphs for a database interface, *IBM J. Res. Dev.* **20**(4) (1976) 336–357.
30. Sowa, J. F., *Conceptual Structures: Information Processing in Mind and Machine* (Addison-Wesley, Reading, MA, 1984).
31. Sowa, J. F., Semantic networks, in: S. C. Shapiro (Ed.), *Encyclopedia of Artificial Intelligence* (Wiley, New York, 1987) 1011–1024.
32. Sowa, J. F., Foo, N. and Rao, A. S. (Eds.), *Conceptual Graphs for Knowledge Systems*, to appear.
33. Sowa, J. F. and Way, E. C., Implementing a semantic interpreter for conceptual graphs, *IBM J. Res. Dev.* **30**(1) (1986) 57–69.
34. Tesnière, L., *Eléments de Syntaxe Structurale* (Librairie C. Klincksieck, Paris, 2nd ed., 1965).
35. Walker, A., McCord, M., Sowa, J. F. and Wilson, W., *Knowledge Systems and Prolog* (Addison-Wesley, Reading, MA, 1987).
36. Way, E. C., Dynamic type hierarchies: An approach to knowledge representation through metaphor, Ph.D. Dissertation, Dept. of Systems Science, State Univ. of New York at Binghamton, 1987.