

Book Reviews

Knowledge-Based Systems

Simon Y. Blackwell

The two-volume set entitled Knowledge-Based Systems (Volume 1, *Knowledge Acquisition for Knowledge-Based Systems*, 355 pp., and Volume 2, *Knowledge Acquisition Tools for Expert Systems*, 343 pp., Academic Press, San Diego, California, 1988), edited by B. R. Gaines and J. H. Boose, is an excellent collection of papers useful to both commercial practitioners of knowledge-based-systems development and research-oriented scientists at specialized centers or academic institutions. The set is the result of a call for papers to support the first American Association for Artificial Intelligence Knowledge Acquisition for Knowledge-Based Systems Workshop, held 3–7 November 1986 in Banff, Canada. Although the conference was held three years ago, these volumes are still timely and sorely needed. Few books dedicated to knowledge acquisition exist.

The first volume, *Knowledge Acquisition for Knowledge-Based Systems*, begins with a paper whose title sounds appropriate: "An Overview of Knowledge Acquisition and Transfer" by the editor B. R. Gaines. The first three pages of the paper are clear and concise. Unfortunately, it then evolves into a rambling philosophical discussion of humanity as a distributed anticipatory system. Obscured by this discourse are several useful observations about the knowledge-acquisition process, for example, that knowledge can be transferred through cultural, verbal, technical, metaphorical, mathematical, and transcendental means.

The next nine papers in the volume present a variety of knowledge-acquisition techniques and training methods. All are good, the most notable being "Generic Tasks for Knowledge-Based Reasoning: The Right Level of Abstraction for Knowledge Acquisition" by T. Bylander and B. Chandrasekaran. Chandrasekaran's papers are usually illuminating, and this one does not fail: He and Bylander re-examine such traditional beliefs as knowledge

should be uniformly represented and controlled and the knowledge base should be separated from the inference engine.

The final 10 papers in volume 1 discuss generalized learning and rule-induction techniques. They are interesting and informative, particularly "Generalization and Noise" by Y. Kodratoff and M. Manago, which discusses symbolic and numeric rule induction. Most rule-induction techniques focus on the use of examples and numeric analysis such as repertory grids. Kodratoff's and Manago's exploration of how the two complement each other is refreshing. Because of their technical nature and the amount of work it would take to put their content to use, most of the papers in this section of the volume are more appropriate for a specialized or research-oriented group.

For those just getting involved in knowledge-based-systems development, *Knowledge Acquisition Tools for Expert Systems* is the more useful volume. In addition to discussing the tools themselves, most of the papers contain details of the knowledge-acquisition techniques that are automated, thus providing much of the same information which is available in the first volume. As an added benefit, they also often discuss the underlying architectures for solving domain-specific problems. For instance, the details of the medical diagnostic architecture laid out in "Design for Acquisition: Principles of Knowledge System Design to Facilitate Knowledge Acquisition" by T. R. Gruber and P. R. Cohen are almost as useful as the discussion of how to build a knowledge-acquisition system.

Volume 2 is particularly germane given the rise in commercial interest about automated knowledge acquisition following this year's introduction of Neuron Data's NEXTRA™ product and last year's introduction of Test Bench by Texas Instruments. Test Bench is actually discussed in "A Mixed-Initiative Workbench for Knowledge Acquisition" by G. S. Kahn, E.

H. Breaux, P. De Klerk, and R. L. Joseph. This volume provides the background necessary to evaluate knowledge-acquisition tools such as NEXTRA, Test Bench, and AutoIntelligence (IntelligenceWare). The vendors of knowledge-based-systems development tools, for example, Inference, IntelliCorp, Aion, AI Corp., and IBM, would do well to pay heed to these books because they point the way to removing the knowledge bottleneck from knowledge-based-systems development.

Overall, the papers in both volumes are comprehensive and well integrated, a sometimes difficult state to achieve when compiling a collection of papers resulting from a small conference. The collection is comparable to Anna Hart's *Knowledge Acquisition for Expert Systems* (McGraw-Hill, 1986), but it is broader in scope and not as structured.

The arrangement of the papers is marred only by an overly brief index. Few readers can be expected to read a collection from beginning to end, and a better index would facilitate more enlightened use. Less important—but nevertheless distracting—is the large number of typographical errors in both volumes.

In conclusion, the set is recommended for both the commercial and research knowledge-based-systems practitioner. Reading the volumes in reverse order might be more useful to the commercial developer given the extra information available in volume 2.

Neurocomputing: Foundations of Research

Terry Rooker

Neurocomputing: Foundations of Research (The MIT Press, Cambridge, Massachusetts, 1988, 729 pp., \$55.00), edited by James A. Anderson and Edward Rosenfeld, is a collection of 43 seminal papers and book sections, each prefaced by several pages of introductory comments. The book