AI Magazine Volume 21 Number 4 (2000) (© AAAI)

Review of *Intelligent Systems for Engineering: A Knowledge-Based Approach*

B. Chandrasekaran

AI has many faces: Its philosophical face is turned to ancient problems about the relation between mind and body, its scientific face to providing a tool for understanding cognition in humans and animals, and its mathematical face to formulating and analyzing classes of algorithms that appear to be effective in providing computers with some of the functions of intelligent activity. For several decades, there has been another face to the field, a technological one that provides tools for solving practical problems in various domains. Of these domains, engineering and medicine have had the closest interaction with AI. AI ideas in representation and reasoning have been especially relevant to diagnosis and corrective action planning in both domains, and in engineering, design has been another area of very fruitful collaboration.

These disciplines have not been mere consumers of AI ideas and technology, however. They have had a deep effect on the theoretical side of AI. Complex real-world reasoning tasks, such as those in engineering and medicine, bring out the inadequacy of purely theoretical conceptions about the nature of intelligence. The enormous amounts and types of knowledge often required for carrying out practical reasoning, and the variety of inference techniques that are displayed by practitioners, and the need to arrive at conclusions rapidly—these issues that once were considered mere "engineering" issues turned out to have nontrivial theoretical content for AI. Focus on classes of tasks, such as diagnosis and design, resulted in identifying characteristic common representational primitives and reasoning strategies. Another area of significant progress, instigated by the focus on problems in engineering and medicine, has been in the representation and use of causal knowledge, and the associated body of ideas for simulation. There is more, but the preceding list gives some idea about how concerns with application bring advances in theory, as has happened earlier in mathematics and physics.

Many academic researchers have found that AI often elicits greater interest from fellow academics in engineering departments—many computer science departments are housed in colleges of engineering—than even other faculty in computer science. AI courses at The Ohio State University, for example, regularly have many enrollees from sister departments of engineering, most of them graduate students. Their or their advisers' interest is typically piqued by a talk with a colleague who has used some AI techniques, or perhaps by a paper that they read in their professional journals about a solution to some problem of interest using AI representational ideas. An initial interest of this sort often blooms into a long-term research program. I know teams of engineering faculty, in chemical, mechanical, and civil engineering and in architecture, who have been carrying on basic and applied research in engineering problem solving using AI ideas and, in fact, have made important contributions that are of interest to the general AI community, not just engineering practitioners. The author of the book under review is, in fact, one such engineering researcher—he built interesting systems using AI ideas for his Ph.D. research in civil engineering at Carnegie Mellon University and then continued investigating issues in representation and reasoning as part of his research career for the last decade and a half.

Intelligent Systems for Engineering: A Knowledge-Based Approach, Ram D. Sriram, Springer Verlag, 804 pp., 1997, ISBN 3-540-76128-4.

However, the engineers, as is their wont, have their own take and emphasis on AI issues. For one thing, they tend to get more excited by matters that have a likelihood of making a difference in their work than matters that are of pure computing or mathematical interest. Teaching engineering students interested in AI, especially when they are taking courses along with computer science students, presents a bit of a challenge because of the difference in background and interest. Also, when ideas are presented somewhat abstractly, the engineering students might need to do extra work in seeing how they might be applied to the problems of interest in engineering. It would thus be great if there were a book that presents, fairly comprehensively, a full range of AI ideas and techniques but in a form readily consumable by engineering students or those interested in engineering applications; uses engineering examples all through to motivate the techniques; and respects the differences in interest and emphasis between typical computer science and engineering students.

Well, we now do have such a book. The book under review is a very comprehensive text that covers an enormous amount of material on the application of AI to engineering problems. In spite of the subtitle "A Knowledge-Based Approach," the book in fact covers even neural net and genetic algorithm approaches, which are usually considered as alternatives to so-called symbolic AI, that is, approaches that can largely be viewed in the representation-reasoning-search paradigm of traditional AI. However, the focus is on approaches that involve knowledge representation, reasoning, and search. In keeping with the intended audience's interests, the book avoids theorems and proofs about the algorithms and representations but focuses on a clear exposition of both basic and advanced ideas, all explained with liberal use of engineering examples.

After an introduction to knowledge systems and the potential for their application to tasks of interest to engineers, the book gives a clear exposition of a variety of basic search techniques and moves on to presenting a range of knowledge representation techniques and inference methods. A fairly large section is devoted to qualitative reasoning, which is appropriate because qualitative reasoning has been a major contributor to the solution of engineering-related problems. Case- based reasoning is discussed briefly. In keeping with the focus on engineering, there is a discussion on management of knowledge-based–system projects.

There is a section on system-building tools—many of them at one point commercially available, but I don't know how many of them are today available in the form described—that has a dated look to it if one views the tools as currently live options. However, the section is interesting as an historical account of the evolution of tools and also as a presentation of the kinds of features that facilitate system building.

Given that the book is more than 750 pages, I don't claim to have read the entire book, but I sampled many sections, especially dealing with topics that I know about or even have contributed to. I especially found the long sections on qualitative modeling and simulation very useful, particularly from the viewpoint of teaching these ideas to engineering students.

In summary, I think this is a valuable book for instructors teaching AI to engineering students and as a handbook and reference source for engineers who use AI in their applications or research.

B. Chandrasekaran is the *AI Magazine* Book Review Editor. He is professor emeritus at The Ohio State University.