

# Book Reviews

---

## Artificial Intelligence: A Knowledge-Based Approach

Joseph Montanarella

To be considered exceptional, a textbook must satisfy three basic requirements. First, it must be authoritative, written by one with a broad range of experience in, and knowledge of, a subject. Second, it must effectively communicate to the reader, in the same manner in which a course instructor must be capable of imparting knowledge to students in a classroom. Third, it must stimulate the reader into thinking more deeply about the subject and into viewing it from fresh perspectives.

In *Artificial Intelligence: A Knowledge-Based Approach* (Boyd & Fraser, Boston, 740 pp., \$48.95), author Morris W. Firebaugh has succeeded in meeting each of these requirements. To his considerable credit, he has created a well-written and comprehensive introduction to AI. It covers all the basics and then some, leaving the reader with a well-rounded understanding of the field. In addition, the book makes for enjoyable reading, which is indeed a rare and commendable feat for a scientific text.

As a graduate student, my preference was for textbooks that began by acclimating me to the subject matter. The first three chapters of this book accomplish precisely this goal. Although these chapters are perhaps a bit overambitious in terms of the volume of material covered, the reader is nonetheless left with an overall feel for what AI is, where it came from, and where it might be headed. Through the introduction of the major features of Lisp, Prolog, Smalltalk, and Pop-11, one gains an appreciation for why these special programming languages are needed for AI applications. The presentation of search techniques is handled particularly well, beginning with descriptions of several problem situations requiring a search strategy and

followed by a discussion of the search techniques available to solve them.

Game playing, the subject of chapter 4, is an interesting topic in its own right and a good choice to illustrate the concept of search heuristics. Having already been convinced of the necessity of limiting the search space, you are clearly shown how the application of appropriate heuristics can effectively limit the search. You begin to understand the role of AI as a separate discipline of computer science by understanding the distinction between problem domains where strictly algorithmic solutions suffice and those where AI approaches are required.

Firebaugh's attempt to describe automated reasoning in Chapter 5 falls victim to an abundance of terminology, laws, and syntax rules with inadequate definition and explanation. The treatment assumes a prior knowledge of logic and propositional and predicate calculus. The chapter on problem solving, however, is Firebaugh at his best. The reading is interesting, and the text flows easily. All the topics in the previous two chapters are neatly tied together, and are well illustrated with good examples and figures. During a fascinating discussion of the general problem solver, the importance of separating the general problem-solving capability of a program from the specific domain knowledge becomes apparent. A brief survey of other problem solvers nicely rounds out the chapter.

The chapter on natural language processing starts with an engaging discussion of communications between humans and computers using the English language. The separate problems of syntax and semantics are discussed, with demonstrations of how the rules of grammar help to define syntax and how the inclusion of knowledge is essential in dealing with semantics. Excellent use is made of Prolog and Lisp code segments as examples, and new

concepts are introduced at a perfect pace and are fully defined. Chapter 8, however, is afflicted with a malaise similar to that of Chapter 5 and, as such, provides little further insight. An example is the discussion of transforming the surface representation of a communication to a deep representation of its meaning. Here, reference is made to a meaningless figure consisting of two boxes, an arrow connecting them, and an uninformative caption, along with sample listings of nothing more than input and output. I was left with no idea of how the transformation actually occurs.

Chapters 9 through 13 constitute the heart of the book and are arguably the best chapters in the text. They concern the power of knowledge, how it is best represented, and how it is most effectively applied. The order of presentation is superb, with each subsequent chapter building on the previous one. We are taken on a journey through the classic knowledge representation techniques, illustrated by several benchmark expert system development efforts. In a move from specifics to generalities, the properties of knowledge-based systems, features of human problem solvers, and architectural principles are explored. A most impressive feature—the construction of a simple rule-based system using a commercially available development tool—comes next. The presentation is in an interactive mode, with the printed page reflecting exactly what you would see if seated at a terminal. This approach is as close as one can get to hands-on learning from a textbook! Supplemented with clear diagrams and examples, the material in each chapter flows smoothly into the next. Although introductory in nature, these chapters nonetheless manage to provide a thorough understanding without overwhelming the reader.

The book also includes chapters on pattern recognition, computer vision, and robotics. Although these areas are unquestionably treated thorough-

ly and informatively, there is simply too much theory and too many formulas to be comfortably digested. Because of the emphasis on the specifics in each field rather than the general principles involved, the relationship between AI and the subjects under discussion is unclear. Some topics, such as Teitelman's algorithm for recognizing handwritten characters, were extremely difficult to decipher. The appropriateness of including such in-depth treatments in an introductory text on AI is questionable. Thankfully, Firebaugh reverts to his more characteristic style with the subsequent chapter on machine learning. Here is a highly focused discussion; the concepts, applications, and relative merits of various machine learning techniques and their relationship to AI are neatly presented.

The reader is left to ponder the future direction of AI in the closing chapter. Current limitations and the potential of parallel processing to alleviate some of these limitations lead naturally into a detailed discussion of neural networks. A great deal of information is available here for the uninitiated. Maligned early on, as was the entire field of AI for promising more than it could deliver, neural networks are staging a comeback, and solid evidence is presented for the continuation and expansion of current research efforts.

The book also contains two appendixes. Appendix A, written by Yong Y. Auh, is a fine tutorial on the Lisp programming language. Appendix B contains a summary of commercially available AI systems, including expert system shells, and can be helpful to one seeking tools for a specific application.

One disturbing aspect of this text is that it is littered with poor grammar and typographical errors. The problems include incorrect tense; plural-singular mismatches; improper punctuation; poor sentence structure; and wrong, duplicate, or missing words. There are also more than a few instances of inappropriate splitting of paragraphs and sections between adjacent pages and of ill-advised placement of illustrations. These lapses are unfortunate because they detract from the reading pleasure and unnecessarily degrade an otherwise well-written book. To alleviate these problems, a thorough proofreading is a must prior to pub-

lishing the next edition.

A textbook is intended to serve as one of the three main components of the learning environment, along with classroom instruction and homework assignments. Although a textbook is primarily used to supplement the instruction and to refer to when completing assignments, this text goes beyond these minimum requirements. It is a valuable learning tool and would be highly informative even if read without benefit of the other two components. The major flaw in Firebaugh's book is the author's tendency at times to introduce too many new concepts with too little explanation.

Firebaugh states in his introduction that his objectives were "to introduce students to the major ideas of Artificial Intelligence [and] to provide the student with at least a conceptual introduction to all major areas of AI and a working knowledge of many of the more practical and applied areas." I can attest to the fact that these objectives were indeed achieved.

*Joseph Montanarella obtained his Master's in computer science from the Rochester Institute of Technology in 1989 and is currently employed at IBM in the Research Triangle Park in North Carolina as an associate programmer.*

## The Media Lab

*Lee S. Brownston*

Stewart Brand, of *Whole Earth Catalog* fame, is a technology enthusiast. In 1986, he spent three months in the fantasyland of his choice, MIT's Media Laboratory (formerly the Architecture Machine Group). In his latest book, *The Media Lab: Inventing the Future at MIT* (Viking/Penguin, New York, 1988, 285 pp., \$10, ISBN 0-14-009701-5), he tells the world what he found. Brand's enthusiasm is infectious: His report would make any working programmer ache to be part of this heady mix of visionaries, hackers, artists, scientists, and entrepreneurs, where flakiness is a positive value, and freedom and money are seemingly without bounds. However, this book is not for technically sophisticated readers alone; it is the kind of book you can give your grandmother to illustrate the existential pleasures of

engineering, in Samuel Florman's phrase.

Brand follows the familiar format of much science reporting, carefully describing and sketching the personalities of the Media Lab stars as hooks to describe their work. His style is breezy and personal, with a sly, sometimes tart wit that somewhat offsets his awe.

In the first of the book's two parts, *The World of the Media Lab*, which composes roughly two-thirds of the book, Brand's personality profiles are skillfully woven with historical and technical background information, anecdotes, and descriptions of the technical achievements; the second part, *The Media Lab of the World*, analyzes the social implications of the technologies being pursued at the Media Lab and elsewhere. There are numerous black-and-white photographs of the people whose work is discussed and a section of color photographs illustrating the work environment and the projects. As a bonus, embossed on the cover is a white-light hologram that illustrates one of the laboratory's successful research projects.

Brand evidently found the Media Lab people to be congenial; they are described so warmly and respectfully that if he didn't truly admire and like everyone whom he wrote about in this book, he did a good job of hiding the fact. Sometimes, the author is so dazzled with lab members that he inserts lightly edited transcripts of his interviews with his interlocutors; these transcripts are not as fascinating to the reader as they evidently were to the author. For example, a long interview with Andrew Lippman is presented almost verbatim to illustrate some of Lippman's ideas concerning conversational discourse. However, the didactic value of performing the transcription must have been greater than the value of reading the result. Granted the interviewees are articulate and present valuable ideas, but a little more digestion would have been appropriate.

Brand's admiration is in marked contrast to the evident distaste of Sherry Turkle: In her book *The Second Self*, she made little attempt to conceal her revulsion at the personalities, tastes, and ideology of many of the same people. Where Turkle stressed the nerdiness and intellectual arrogance of MIT's