

BOOK REVIEWS

Artificial Intelligence for Microcomputers

If you would like to develop an expert system or knowledge-based system on a microcomputer, you might want to read *Artificial Intelligence for Microcomputers* by Mickey Williamson.

This nontechnical book is easy to understand, written for the unsophisticated microcomputer user. The first chapters provide a brief history of artificial intelligence (AI) and an introduction to natural language query systems. They explain what knowledge-based systems and expert systems are and how they work. Discussions are also provided of the two major AI programming languages, Lisp and Prolog, including their strengths and weaknesses. The remainder of the book is devoted to a review of some of the existing AI software products for microcomputers, such as natural language query systems, decision support systems, expert system development shells, and AI programming languages.

I found the chapter on expert system shells to be worth the price of the book. An expert system development shell enables the user to develop a knowledge-based system without having to write programs. The shell provides a way of representing knowledge and a way of reasoning from that knowledge to reach a conclusion. This chapter not only provides criteria for choosing a shell but also describes the features of six different shells. One of these shells, EXSYS, offers a great deal of power and flexibility at a very reasonable price. It uses if/then rules to build the knowledge base. You can add explanations and references to the rules. The rules are written in plain English and do not require any knowledge of programming. You can assign probability factors to the rules. The system can explain what rules it employed to reach a conclusion. When you are developing an application, you can stop anywhere in the program and rerun it to check the validity of the rules to that point. EXSYS has extensive mathematical capability, including the ability to perform trigonometric functions, call external programs, import data from these programs, receive input from automated testing equipment, and reason forward to a conclusion or backward from a goal.

EXSYS is just one example of the type of software reviewed in this book. The book also describes some very interesting natural language query systems that allow you to retrieve information from a database using conversational English; a number of consultation systems that help you reach decisions are covered as well. By the time you finish reading this book, you will know most of the vocabulary used in AI systems; you will know how AI evolved to where it is today and where it is likely to go; and you will know about a number of inexpensive AI products available for microcomputers. This book is an excellent source for someone

who would like to do something in the field of AI without spending a great amount of money.

Artificial Intelligence for Microcomputers: The Guide for Business Decision Makers. Mickey Williamson. (New York: Brady Communications Company, 184 pages, \$17.95).

T. R. Vogel United States Geological Service,
National Mapping Division
Menlo Park, California 94025

Introduction to Artificial Intelligence

Writing about artificial intelligence has turned out to be a considerably more difficult undertaking than many people had suspected. Even writing an introductory text for this rapidly growing field is a challenging task whose difficulty has generally been underestimated. My intention in this review is to look at one of the best attempts to date at supplying such an introductory text to the field of AI and to do so in the spirit of perhaps arriving at some useful standards against which such an attempt might be gauged.

A major shortcoming of many of the introductory AI volumes is the use of material which is so out of date that the book projects a false impression of the field as currently practiced, making the book's usefulness highly questionable. This point alone places the text by Eugene Charniak and Drew McDermott in a special category because its focus is primarily on problems currently being addressed in the AI field. Some aspects of the material presented might even be appearing for the first time in this volume.

The authors indicate in the preface that they intend the book for use at the undergraduate level. This aspect, however, should not be considered a limitation of the book's level of sophistication and depth. As a matter of fact, this book seems quite suitable for use by people who already have some background in AI.

The text is organized into eleven chapters. The first chapter provides a brief introduction to AI and what it is and to various approaches to knowledge representation, particularly the predicate calculus notation that is used throughout the book for developing examples. Chapter 2 contains a succinct introduction to Lisp, which uses a dialect similar to that of Franz Lisp. The rest of the book is devoted to particular AI application areas and methodologies, such as vision, natural language, logic programming, search, expert systems, planning, and learning.

This book is an extremely good companion volume to *Artificial Intelligence Programming* (Lawrence Erlbaum Associates, 1980), an earlier volume on which the authors collaborated with Christopher Riesbeck. As opposed to the

earlier text, the coded examples used here are sketchy; as a result, they are often difficult to follow or understand. But here we run up against a dilemma that all introductory AI books must face: whether to include complete listings of AI programs. Including the listings seems to take such a book out of the introductory category, but incomplete listings are seldom useful, and no listings at all makes the book appear too superficial. The solution seems to be one this book utilizes to some degree: make the complete listings available in an appendix for those who wish to make use of them.

The authors provide a diplomatic overview of AI that includes many contrasting and competitive approaches which are seldom covered together in the same volume with comparable sympathy and detail—an admirable achievement. The result is a broader and more enriched view of the field than is usually conveyed. The chapter on vision systems divides the problem of interpreting visual images into the two phases of early processing and scene representation and recognition, providing lengthy sections on both problem areas. The chapter on natural language parsing illustrates how syntactic parsing can be augmented with semantic parsing. It presents the outline of a Lisp program that uses ATNs for the syntactic parser. Various high-level functions such as *def-feature*, *fed-net* and *def-semantics* are introduced and used to outline a system. The treatment suffers, though, from an inadequate explanation of these functions, what they do and how they are implemented. Similarly, in the chapter on search, the outlines of search algorithms and Lisp functions are presented, but the reader is left again with the feeling that too much is left undisclosed about these functions to permit an adequate understanding.

The chapter “Logic and Deduction” contains an interesting discussion of the conceptual dependency approach to natural language understanding that formulates CD representation in predicate calculus. Other interesting topics in this chapter include nonmonotonic reasoning and modal and intentional logic. Perhaps the most intriguing chapter is “Memory Organization and Deduction,” which touches upon the topics of frame-based representation and deductive retrieval and introduces the time-order representation approaches of temporal system analysis and time map management.

One of the most uneven chapters in the book is that on expert systems. Although the discussions of uncertainty and Bayesian probability are extremely clear and illuminating, the unique nature of production systems, the rule-based programming paradigm, and the various issues involved in the use and development of expert system technology are not elaborated. The final chapter, which is devoted to learning research, provides a very broad survey of nearly all the main directions in this important field of AI research. The topics range from rule induction to learning by analogy to failure-based learning; the chapter even includes a section on Lenat’s discovery program Erisko.

In my opinion, the best chapter in the book is “Managing Plans of Action.” Here, the authors present an unique approach to temporal planning which has been used at Yale in the Forbin project, an experimental robot planner that attempts to optimally plan multiple tasks, each of which is under the constraint of a deadline. A nice touch is the outline for a computer program that plays the game of Risk.

Abduction is an important new term introduced in the book. Abduction in this sense means a form of inference that moves you in the opposite direction from deduction. Rather than providing correct logical conclusions, abduction is a means of tracing a chain of inferences backward for the purpose of explaining conclusions.

An interesting point the book makes is that “backtracking is just another name for depth-first search.” The argument goes like this: Depth-first search involves maintaining a queue of states that have been generated but are still unexplored. To proceed in a depth-first manner, then, means to manage “last in-first out” like a stack. This approach is the same used in backtracking: the nodes remain at lesser depths, still on the stack to be explored if the search reaches a dead end. Such a clear explanation allows the student to see the underlying simplicity behind the complexities of AI search algorithms.

On the whole, the book is filled with valuable and occasionally brilliant insights into solving AI problems. At times, however, statements are made categorically, as though they were self-evident, but not enough explanation is offered to allow the reader to fully appreciate the insight. One of the book’s strongest merits is that it does not ignore the issues of practical implementation and efficiency as trivial details but gives them serious consideration on nearly every page.

Without question, this text is the best general introduction to AI that I have seen. However, although masterful and superb, it still contains inadequacies. This statement is more a testament to how remarkably difficult it is to write an adequate introductory AI text than it is a criticism of the job done by the authors. Because there really are no single volumes yet that provide a satisfactory introduction to AI, the best way to approach the problem of selecting text material for an introductory AI course seems to be use a book such as this which comes close to the mark in several ways, and supplement the weaknesses of the text with additional materials. As is true with many aspects of AI today, the very best is not always enough to satisfy all our needs and requirements.

Introduction to Artificial Intelligence. Eugene Charniak and Drew McDermott (Reading, Massachusetts: Addison-Wesley, 701 pages, \$25.00).

Ernest R. Tello
Integral Systems
Santa Cruz, California