J. Indian Inst. Sci., Mar.–Apr. 1990, 70, 187–196.

BOOK REVIEWS

Neural networks and natural intelligence edited by Stephen Grossberg. The MIT Press, 55, Hayward Street, Cambridge, Massachusetts, 02142, USA, 1988, pp. 637, $35. Indian orders to: Affiliated East-West Press, 25, Dr. Muniappa Road, Kilpauk, Madras 600010.

This book is a collection of 12 research articles written by Grossberg and his associates at the Center for Adaptive Systems, Boston University. Professor Grossberg's group has been studying the design principles and mechanisms needed for a biological or technological system to exhibit intelligent behaviour. This book, which brings together the important results of this investigation, mainly addresses the problems involved in realising real-time goal-directed adaptive behaviour. The topics discussed include visual processing (form, color and brightness perception and stereopsis), adaptive categorisation of patterns, associative search, attention, competitive learning, planned motor control and decision making under risk.

The main theoretical idea presented in the book is adaptive resonance. This is a self-organising, self-stabilising neural network model for pattern recognition. There are two conflicting design requirements when one is concerned with an organism or system that can learn to categorise input sensory patterns through its interactions with a complex environment. The system should be stable in the sense that an input pattern, which is only a small variation of an already classified pattern, does not cause change in the behavioural response. At the same time, the system should be adaptive or plastic so that 'novel' changes in the sensory pattern will make the system evolve to include a new categorisation and the needed behavioural output. This conflict is referred to as the stability-plasticity dilemma. In classical supervised pattern recognition, one assumes the availability of a teaching signal after each classification by the system and the various categories to be learnt are prefixed. In unsupervised methods such as clustering, the type of data regularities that the system is sensitive to are wired-in *a priori*. Adaptive resonance theory (ART) incorporates mechanisms such as attention, multiple interacting memories with associative search, and a modifiable vigilance system to overcome the stability-plasticity dilemma. Here the vigilance system triggers learning of the recognition code for a new category when there is 'sufficient' misfit between the new pattern and the learnt recognition codes. This vigilance itself is sensitive to the complexity of patterns and also the history of learning so far. Further, unlike the classical algorithms, here one is interested in designing systems where learning comes out as an emergent behaviour of a large number of interacting elements and subsystems. While ART is not a complete theory of learning and pattern recognition, it contains many ideas which are worth pursuing further.

While the ability to learn to categorise is an important component of intelligence, a complete system should also include abilities such as: perceptual processing of raw sensory input into pattern input for the recognition system, decision-making ability so as to decide on the behaviour based on the perceptual input and the current goals, and ability to coordinate motor activity to achieve the calculated behaviour. This book contains articles that discuss possible neural architectures for all these problems. With such a holistic approach, Grossberg and his colleagues are able to suggest important models for understanding human cognition and also throw some light on the vast amount of data collected on biological systems.

The first four chapters of this book address the problem of preattentive vision. These discuss neural architectures for segmentation and filling-in features, characterising surfaces, stereopsis, color perception, etc., to obtain a preattentive representation of three-dimensional form. Chapters 5–8 discuss adaptive pattern recognition, ART, and competitive learning. Chapters 9–11 include architectures for rapidly focussing attention and computing needed behaviour for satisfaction of internal goals. Chapter 12 discusses a possible architecture for motor control and robotics.

The book is very hard to read and it requires considerable effort to understand the material presented. A good part of this difficulty is due to the nature of the subject matter. As Grossberg indicates in his preface, all the authors have advanced training in at least three of the four areas: computer science, mathematics, psychology and neurobiology (the 3/4 rule). So it is only to be expected that people who are familiar with only one or two of these areas will find the material hard. Artificial intelligence and the recent field of neural networks are interdisciplinary by the very nature of questions they address. Hence people who are seriously interested in these areas will find the book very useful. The papers presented here contain a wealth of data on biological systems and also an extensive bibliography. Many theoretical ideas presented here are worth pursuing. For people who have not come across this type of material earlier, Grossberg's article in *IEEE Computer*, March 1988, would be a good starting point. In this book, perhaps the best papers to start with would be chapters 6 and 7. The chapters on vision would be of interest only to people with prior experience in computer vision. The remaining chapters would be of interest to computer scientists with a psychology background. On the whole, the book is a useful collection of papers that serves as reference material for researchers interested in neural modelling of intelligent systems.

Department of Electrical Engineering
Indian Institute of Science
Bangalore 560012.

P. S. SASTRY

The artificial intelligence debate: false starts, real foundation edited by Stephen R. Graubard. The MIT Press, 55, Hayward Street, Cambridge, Massachusetts, 02142, USA, 1988, pp. 328, $9.95. Indian orders to: Affiliated East-West Press, 25, Dr. Muniappa Road, Kilpauk, Madras 600010.

The field of artificial intelligence (AI) has, by the very nature of its subject matter, always attracted the attention of many disciplines other than computer science. Since understanding the functioning of mind and endowing computers with cognitive capabilities is as much of an engineering problem as a philosophical one, the history of AI is replete with philosophical arguments and counter arguments. With the recent resurgence of artificial neural networks (or the so-called parallel distributed processing models), there are now at least two schools of thought within computer science as to the right direction towards the ultimate goal of AI—one based on the traditional symbol processing approach with logical reductionism as its backbone and the other based on the connectionist paradigm which may be loosely termed as Gestalt Neuroscience. The debate as to which is more correct is going to be a central theme of AI in the years to come. The present book brings together highly thought-provoking essays on this issue from a wide variety of viewpoints with computer scientists, philosophers, biologists all joining the debate.

In the first article, aptly titled "One AI or many?", Papert sets the stage for the discussion to follow in the rest of the book. The next paper, by Hubert and Stuart Dreyfus, is a thorough discussion of the philosophical assumptions involved in the two contending approaches to AI. This essay also gives a fairly good historical account of the development of the two schools of thought.

The next two papers by Sokolowski and McCorduck represent a spectrum of viewpoints mainly from the symbol-processing tradition. They try to look beyond the current capabilities of AI programs and attempt to characterise what a physical system with significant symbol-processing capabilities could mean.

The paper by Cowan and Sharp presents a fairly comprehensive introduction to artificial neural networks and their relevance as models of intelligent systems. For people trained mostly in traditional AI, this can serve as a tutorial to better appreciate the discussion in other papers. The next two papers discuss the limitations of connectionist models *vis-a-vis* real neural systems. Schwartz describes how complex real neurons are and suggests that the models should capture some amount of the biological diversity of neurons rather than settle for uniform mathematical models. In a more forthright criticism, Reeke and Edelman point out that the current research in AI is mostly irrelevant to understanding human mind because it largely ignores biological information. This paper also gives a glimpse of neuronal group selection theory, one of the current biological models for understanding brain functioning. What emerges out of these two papers is that there should be more interaction between AI and neurobiology much like that between theoretical and experimental physics.

The next three papers by Hillis, Waltz, and Hurlbert and Poggio discuss how the two approaches should be combined for realising the ultimate goal of understanding and replicating intelligence. Hillis suggests that connectionist systems characterised by emergent behaviour and traditional AI characterised by symbolic processing should be viewed as complementing each other. "Choosing between emergence and symbolic computation in the study of intelligence is like choosing between metabolism and genetic replication in the study of life. Just as the metabolic systems provide a substrate in which the genetic system can work, so an emergent system may provide a substrate in which the symbolic system can operate". Hurlbert and Poggio also voice a similar viewpoint, taking machine vision as an example where this synthesis already exists. They also suggest that Marr's argument for studying any intelligent system from three different levels, contains a suggestion of how this synthesis can be attempted.

The two papers by Putnam and Dennett represent the two extreme viewpoints philosophers have on AI. Putnam dismisses AI as philosophically irrelevant while granting that it may have some engineering significance. Dennett sharply disagrees with Putnam and tries to project AI programs as more accurate philosophical thought experiments if only because they offer a possibility of empirical falsifiability of philosophical hypotheses. The last paper by McCarthy is somewhat disjointed with the rest of the book. McCarthy traces the development of logic formalisms for AI and simply acknowledges the implicit philosophical assumptions.

This book is originally published as special issue on AI in the Winter 1988 number of *Daedalus*. The book contains a wide spectrum of viewpoints and it is naturally impossible to agree with all of them. While no paper contains any new results or theories, each paper contains ideas which are worth pondering over. The book should be read by all AI researchers who are more interested in understanding human mind than in building commercially viable robots. The book is easy to read for anybody with some knowledge of the history and techniques of AI. Thus all AI researchers in academic community would enjoy this book.

Department of Electrical Engineering
Indian Institute of Science
Bangalore 560 012.

P. S. SASTRY

What every engineer should know about artificial intelligence by William A. Taylor. The MIT Press, 55, Hayward Street, Cambridge, Massachusetts, 02142, USA, 1988, pp. 331, $25. Indian orders to: Affiliated East-West Press Pvt. Ltd, 25, Dr. Muniappa Road, Kilpauk, Madras 600 010.

The field of artificial intelligence (AI) often conjures up visions of humanoid robots taking overall traditional chores of man. There are many intense philosophical battles on whether machines can ultimately be endowed with intelligence and even if they can be, whether what AI does is the right step in this direction. If making computers behave like people is the only aim of AI then the endeavour should largely be classified as failure because 'Silicon People' appear to be far away in future. But there is another more practical side to AI. Research in AI has resulted in many techniques and insights into how computers can be better utilized in many different fields. In this book Taylor takes the view that 'for engineering purposes artificial intelligence is a few software ideas that work well enough for commercial use' and goes on to explain in a lucid and highly readable manner how techniques of AI can be utilized by engineers wanting to harness power of computers for solving different kinds of problems. Due to his background as a consultant, the author is in an ideal position to explain the commercial relevance of AI techniques.

The book starts with an engaging account of what AI is all about and what are some of the offshoots of AI research. The discussion is at a popular level and no technical background is needed for enjoying the material. Then, in Chapter 4, the author gives more than ten specific case studies of engineering problems where AI techniques were found useful. This concludes what can be called the first part of the book.

In the second part of the book the author explains some of the basic concepts of AI which have relevance to developing software systems. There are simple introductions to LISP and PROLOG, chapters on programming styles which explain the ideas behind object-oriented programming, function-based programming and rule-based programming. The presentation is suitable for a working engineer who has no background in AI but is familiar with computers and programming. Any book of this kind has to necessarily talk about expert systems if only because they are the most visible success of AI. The book contains one chapter on expert systems where the basic principles are explained and another on the important aspect of writing friendly user interfaces to expert systems. There is a separate chapter on logic programming which is possibly the best written chapter of this book.

There is a chapter on future trends in AI where the author puts down his ideas about some non-conventional use of AI techniques. The final chapter deals with Japanese fifth generation project and the author's ideas about how American enterpreneurs can fight back the 'Japanese onslaught'. This might be of interest to American audience particularly because the author grew up in Japan.

The book is mainly meant for working engineers and managers of software projects. It contains many tips about what kinds of project are suitable for AI techniques and how to convince top management about developing such systems. While the explanation of basic ideas is easily understandable by a novice, the book itself is written in a sales pitch. The author's intention seems to be to vigorously sell the idea of AI techniques. Thus at some places the narrative becomes an advertisement for AI and the style becomes cliche-ridden. But on the whole, the book should be useful for managers who can learn about AI techniques at an elementary level and make decisions about what kind of software projects have better potential for future.

Though the book is highly readable and contains many little anecdotes and jokes, for students of computer science and AI there is nothing much to learn. For the academic community, possibly the

most interesting part would be Chapter 4 where many specific industrial applications are described. This can give a feel for the kind of problems that exist in the real world.

Department of Electrical Engineering
Indian Institute of Science
Bangalore 560 012.

P. S. SASTRY

Artificial intelligence in psychology by Margaret A. Boden. The MIT Press, 55, Hayward Street, Cambridge, Mass. 02142, USA, 1989, pp. 188, $20. Indian orders to: Affiliated East-West Press Pvt. Ltd., 25, Dr. Muniappa Road, Kilpauk, Madras 600 010.

This book contains eight interdisciplinary essays on the general theme, relevance of artificial intelligence (AI) to psychology.

Computational psychologists consider psychology as the study of various computational processes whereby mental representations are constructed, organized, interpreted or transformed. Most computational psychologists further assume that these processes need be understood only at the level of symbol-manipulations defined in terms of formal syntactic rules. This philosophical assumption is known to AI scientists as the physical symbol system hypothesis. As Boden points out in her introduction, this assumption can be seen even in Ada Lovelace's writings on Charles Babbage's analytical engine.

However, during the past few years a new type of computational models are gaining popularity as a possible alternative to traditional symbol-processing approach. These models are variously known as neural networks, parallel distributed processing (PDP) models or connectionist models. In a recent special issue on AI of the journal *Daedalus*, Seymour Pappert wrote an article titled "One AI or many?" This title aptly sums up the current response to connectionist models. Is connectionism a 'devastating challenge' to conventional AI and computational psychology, as some critics feel? Should it be described as a 'radical departure' from the symbolic paradigm and introducing a 'sub-symbolic' one? Or is it merely a way of implementing the computational ideas of AI on a specific hardware? Should connectionism be viewed only as a model of neural circuitry or can it be viewed as a general information processing model different from the traditional one? Finally, are traditional AI and connectionism fundamentally opposed to each other or can they be usefully combined for a better understanding of intelligence? These are some of the important questions raised by the new developments in connectionist models. Margaret Boden asks all these questions and tries to answer them in so far as they concern theories of psychology.

Chapter 2 of the book contains a concise introduction to computational psychology set in the historical background of changing trends in psychological theories. In the next paper titled "Is computational psychology constructivist", Boden touches upon the old debate of whether perception can be viewed as a computation. This chapter gives a good overview of Marr's theory of low-level vision and discusses Marr's views on the various levels at which a computational task should be understood. The fourth chapter titled "Does artificial intelligence need artificial brains", addresses the question of how relevant or essential is the special architecture of brain for exhibiting intelligence. This chapter contains a good discussion on the possible relationship between traditional AI and connectionism. All these three papers are accessible to a wide range of audience with no special background in psychology. In particular, computer scientists with some background in AI can easily appreciate the discussion presented here.

The remaining five chapters presuppose some knowledge of the various psychological and philosophical issues that are debated in connection with AI. Chapter 5 discusses the classical problem of intentional mental states and whether or not intentionality can be explained in basically mechanistic terms. The next chapter contains a refutation of a famous thought-experiment designed by J. R. Searle for proving that it is futile to use AI concepts in psychology. This so-called Chinese Room problem is originally devised to show that the programs, developed by Roger Schank, which can answer questions based on an English story, cannot be regarded as capable (even in principle) of understanding. This thought-experiment has a long history of refutations and counter refutations. Boden puts forth her own arguments as to why Searle's conclusions are unsubstantiated. Chapter 7, which seems to be slightly out of place in this collection, discusses Piaget's concept of equilibration. Though this has been dismissed by many psychologists as too vague to be useful, Boden suggests that in the light of the new computational methods it is worth taking a second look at the questions that Piaget raises. The next chapter, "Artificial intelligence and biological intelligence", suggests various ways in which AI concepts and methods could be of use to ethologists studying animal intelligence. The final chapter examines whether work in AI and computational psychology has any relevance to theory and practice of education.

This book is very readable and gives a good picture of how some psychologists view the developments in AI. Many questions debated in the book are of philosophical nature and hence cannot have very definite answers. But the discussions are quite illuminating and people interested in understanding human intelligence will find these essays thought-provoking.

Department of Electrical Engineering
Indian Institute of Science
Bangalore 560012.

P. S. SASTRY

An artificial intelligence approach to legal reasoning by Anne von der Lieth Gardner. The MIT Press, 55, Hayward Street, Cambridge, Mass. 02142, USA, 1987, pp. 225, $25. Indian orders to: Affiliated East-West Press Pvt. Ltd., 25, Dr. Muniappa Road, Kilpauk, Madras 600010.

Developments in computing and artificial intelligence (AI) have also started to excite lawyers and other members of the legal profession just as they have excited doctors, engineers and managers. The proceedings of the first International Conference on Artificial Intelligence and Law conducted in May 1987 at Northeastern University, Boston, in cooperation with ACM SIGART provides an overview of the research in this area. Apart from the obvious data base applications of retrieval of legal source materials, applications of AI in the domain of law span over a wide variety of problems such as

—analysis of judicial decision making,
—application of formal logic in the analysis of statutes and other legal documents,
—development of expert systems in specific legal domains such as tax and pension laws,
—modelling legal reasoning from the point of view of a lawyer advising his clients, and
—application of natural language understanding techniques in the interpretation of law.

The present book is a revised version of the author's 1984 Ph.D. dissertation submitted to the Stanford University. The book is thus not a text book in the field, but rather more like a research monograph in the area of AI and legal reasoning. Gardner's survey article in the *Encyclopedia of Artificial Intelligence* describes where her work fits in the area of law applications of AI[1]. Given the two areas as wide apart as law and AI, the fact is that the experts and practitioners in each individual

domain do not fully comprehend both the potentialities and limitations of the problem-solving techniques of AI and the inherent complexity of the legal domain. As a graduate in law from Stanford Law School with many years of legal experience, Gardner is very successful in convincing the computer science reader of the special characteristics of the domain.

The first hurdle for a computer science reader is the legal terminology. He has to understand the many ways in which words like 'case' and 'rule' are used in the law. For example, 'a case' is defined as "a legal dispute with supporting collections of facts and arguments which has been decided by a court"[2]. The word 'case' is used as a synonym for a court's opinion. Similarly a 'rule' may be a legal rule abstracted from several cases. It could indicate a maxim such as "one should not profit from one's misdeeds". Some rules may be statutory and are stated as regulations of a government agency. These are to be kept in mind along with the usage as in AI and Expert systems which uses if-then rules for characterising the 'situation-action' behaviour of a human expert. In this context, facts and rules constitute the knowledge base of the system. The computer science reader has to appreciate the fact that legal language is quite different from the ordinary language. Gardner gives a nice example (p. 43) which may be summarised as follows: "A said he wanted to buy B's car" could be a fact while "A manifested willingness to enter into a bargain with B regarding the latter's car" is not a fact but a legal conclusion. Legal problems also have no clear yes-no answers and conflict, disagreement and argument (on the part of the two lawyers representing both sides of a dispute) are part and parcel of the law. The answer may be more often a "may be" rather than a crisp "yes or no".

Organisation of the book

In the first chapter, the author describes some special features of the legal domain and the relevance of AI techniques for some of the problems in the domain such as those involving the formation of contracts between two parties by offer and acceptance.

The second and third chapters are respectively, titled as Design I: The place of rules and Design II: When the rules run out. In Chapter 2 the discussion starts with an early notion of law called 'mechanical jurisprudence' which essentially embodies the view that law is a set of axioms and legal reasoning is deduction. A subsequent movement in law called 'legal realism' developed as a reaction to mechanical approaches to law. Computer science readers can immediately notice the analogy to the famous 'neat *vs* scruffy' debate in the AI community. The author also addresses the questions of recognising what the legal questions are and how to distinguish the easy from the hard ones. The chapter ends with a discussion on formulating rules for easy questions in the contract law.

The third chapter considers a thoughtful discussion of the 'brittleness' problem of expert systems, that is, what happens when the rule-based reasoning does not work. There is often room for judicial interpretation of a given legal situation, for example, by the use of the word 'reasonable'. Chapter 4 surveys the essential background work on AI applications in law. Chapters 5 and 6 deal with the details of LISP implementation. Gardner's program takes as input a representation of a problem from a law school or bar examination. The output is a data structure similar to a decision tree, in which the decision points correspond to the 'hard questions' that would need to be resolved in order to decide the case. Chapter 7, as the title indicates, deals with the performance evaluation of the program with respect to some test problems. The final chapter is the usual 'suggestions for further work' section of a Ph.D. thesis and contains some insightful observations on 'where to go from here'.

A general discussion of the field of AI and legal reasoning and a discussion of Gardner's book with respect to its utility in legal and computer science education appears in Rissland's work[2] where she describes her experience of using this book in seminar courses to both law and computer science students at Harvard Law School and at the University of Massachusetts, respectively.

In summary, this monograph is a valuable addition to the collection of AI books of any library for providing an insightful introduction to the scope of AI applications in the domain of law. This reviewer has used it in seminars on expert systems applications and also encouraged some students at Indian universities to work in this area. The first such Ph.D. thesis has recently been submitted by Venkateshmurthy[3].

References

1. GARDNER, A. v d. L. — Law applications, in *Artificial intelligence encyclopedia*, Vol. 1, (ed.) S. C. Shapiro, Wiley, 1987, pp. 456–464.

2. RISSLAND, E. — AI and legal reasoning, *AI Mag.*, 1988, **9**, 45–55.

3. VENKATESHMURTHY, M. G. — *Knowledge representation for computer-based legal consultation systems: An experiment with lay-off and retrenchment compensation disputes*, Ph.D. Dissertation, Faculty of Engineering, Anna University, Madras 600025, July 1989.

Department of Computer Science and Automation
Indian Institute of Science
Bangalore 560012.

V. V. S. SARMA

Reasoning about change: time and causation from the standpoint of artificial intelligence by Shoham Yoav. The MIT Press, 55, Hayward Street, Cambridge, MA 02142, USA, 1988, pp. 200, $25. Indian orders to: Affiliated East-West Press, 25, Dr. Muniappa Road, Kilpauk, Madras 600010.

This book describes a method for representing temporal information and reasoning about it. It consists essentially of Shoham's doctoral dissertation and as such it is not intended for the novice. The book is written in a clear and economic style. It demonstrates the author's thorough understanding of the research work carried out in the relevant areas. It assumes familiarity with modal logics, nonmonotonic logics, and temporal reasoning in AI.

The main message of this book has been that mechanical reasoning about the real world can be done in a manner that is both efficient and rigorous.

In the first chapter, two problems that arise from the conflict between the goals of efficiency and rigor in the particular context of the prediction task are identified. These are labelled as the qualification problem and the extended prediction problem.

In Chapter 2, a language for representing temporal information is defined. This is a direct generalization of McDermott's temporal logic. Time point symbols are considered as primitive and time intervals are represented using a pair of time points. Each primitive formula is a pair $\langle i, p \rangle$, where i is an interval symbol and p a primitive propositional symbol. This basic structure is used to construct a categorization of proposition types that is richer and more flexible than the fact/event dichotomy or the property/event/process trichotomy. As a result, the logic rests on a conceptually clean foundation that is at the same time very general.

A new semantic approach to nonmonotonic reasoning is proposed in Chapter 3. It is shown that circumscription, autoepistemic logic, default logic and Bossu and Siegel's nonomonotonic logic are specializations of the proposed logic based on model-theoretic considerations. The general idea is elegant, simple, and consequently, highly appealing. The basic idea employed, a generalized version of

the minimal model concept behind circumscription, is based on ordering the models of a logic by some preference criterion.

The nonmonotonic logic of chronological ignorance (CI) is developed in Chapter 4. It is shown that for causal theories in CI, the set of consequences can be determined in $0(n \log n)$ time, where n is the size of the theory. The author demonstrates that this logic solves the qualification problem.

Chapter 5 is concerned with a slightly more sophisticated use of the logic of chronological ignorance to solve the extended prediction problem. The intuitive concept of potential histories, which are ways the world tends to behave in the absence of interference, was embodied in the class of theories called Inertial theories.

Finally in Chapter 6, it is argued that the preceding work provides a sound basis for a new account of causation. However, there are certain limitations: (i) only primitive propositions can be caused by something, and (ii) it is possible that some counter intuitive statements of causation may be admitted.

To summarize, this is a highly interesting and useful book for researchers in AI. The adoption of a formal framework leads to a clear, principled investigation. The author's solution to the qualification problem and the extended prediction problem using an intuitively appealing nonmonotonic logic based on preferred models is significant. However, the time complexity, given in Chapter 4, of the algorithm for constructing the base sentences is derived without taking into account the number of conjuncts of 0. In practical situations, this may have a significant role to play. Causal theories and potential histories provide a satisfying solution to problems of reasoning about temporal phenomena. The author's style is compact and the reader is expected to have a good background in nonmonotonic logics and modal logics. As such, it may be difficult to use it for classroom teaching. The book is brought out in a nice format. However, there are several typographical errors, scattered throughout the book.

Department of Computer Science and Automation
Indian Institute of Science
Bangalore 560012.

M. NARASIMHA MURTY

Text, context, and hyper text (Writing with and for the computer) edited by Edward Barrett. The MIT Press, 55, Hayward Street, Cambridge, MA 02142, USA, 1988, pp. 368, $35. Indian orders to: Affiliated East-West Press Pvt. Ltd., 25, Dr. Muniappa Road, Kilpauk, Madras 600010.

The dust cover of this book states that this book "presents recent developments in three related and important areas of technical communication: the design of effective documentation, the impact of new technology and research on technical writing, and the training and management of technical writers".

The book is a collection of articles written by technical documentation experts from industry and universities. This book is not another manual on technical writing. Its emphasis is very different. To quote from the foreword of the book "Even computer programs that process only numbers come with a good deal of English. The traditional book is read from the beginning to the end, so the author can write each page knowing that the reader has just finished the previous page. Electronic documentation, however, is not read as much as referred to. The readers can skip around, and usually just look up one item and then another. Often the documentation for a program is available on paper as well as online; the online description comes in pages of a different size than the paper copy, can be opened in different ways, and doesn't have coffee stains to indicate the well-used pages.

The new concept of 'hypertext' and 'hypermedia' describe this new kind of material, which is not so much disorganized as multiply organized: there are many ways to read the material, and many orders in which the pages can be turned. Each reader can have a different experience with the same material and learn different things from it. There is a new opportunity to tailor text simultaneously for many readers. This book tells us a bit about how to do it, and about the psychological and linguistic background to understand why to do it, and what can be achieved."

The book has a thirteen-page introduction by the Editor giving the background on the evolution of thought which led to this book. It is followed by twenty articles organized into three sections: Section I on Artificial intelligence, document processing and hypertext. Section II on Management training, and corporate culture and Section III on Designing on-line information. There are seven articles in Section I, nine in Section II and four in Section III. All articles are up to date and thought-provoking. They are independent articles and range from computer science-based articles such as "Text processing with the START natural language system" by Boris Katz and "Using an object-oriented programming language to create audience-driven hypermedia environments" by Geri Younggren to an article on writing style "How friendly is your writing for readers around the world" by John Kirkman. This article discusses the need to think internationally while writing manuals so that those whose mother tongue is not English can easily understand manuals. I thoroughly enjoyed reading the articles.

This book is an important one and must be read by everyone involved in creating documentation for computer-based systems.

Super Computer Education and Research Centre
Indian Institute of Science
Bangalore 560012.

V. RAJARAMAN

Published by N. M. Malwad, Executive Editor, Journal of the Indian Institute of Science, Bangalore 560 012;
Typeset and printed at Thomson Press (India) Ltd., Faridabad 121 007, Haryana.