



Book Reviews

Perceptrons

Stephen Grossberg

The expanded edition of *Perceptrons* (MIT Press, Cambridge, Mass., 1988, 292 pp., \$12.50) by Marvin L. Minsky and Seymour A. Papert comes at a time of unprecedented interest in the biological and technological modeling of neural networks. One need only refer to the August 1988 issue of *AI Expert* and Philip Chapnick's editorial noting that more than 50 percent of the Fortune 500 companies are investigating neural network technologies. The one-year-old International Neural Network Society (INNS) already has over 3500 members from 38 countries and 49 U.S. states, with members joining at the rate of more than 200 per month. The American Association for Artificial Intelligence was, in fact, a cooperating society at the INNS First Annual Meeting in Boston on 6–10 September 1988. Hardly a week goes by in which a scientific meeting or special journal issue does not feature recent neural network research.

Thus, substantive technical reviews or informed general assessments of the broad sweep of neural network research are most welcome to help interested scientists find their way into this rapidly evolving technology. This fact is especially true because it is the cumulative impact over the past two decades of many scientific discoveries across several scientific and engineering disciplines that has triggered the current explosion of interest in the neural network field.

Unfortunately, neither Minsky nor Papert has participated in any of the many scientific developments that have taken place within the field of neural networks since the first edition of their influential book was published in 1969. In preparing their expanded edition, they have approached the field as outsiders.

Thus, this book contains no new scientific results. Rather, it presents their personal opinions about some recent neural network results in a short prologue (9 pages) and a longer epilogue (34 pages).

In the original edition of *Perceptrons*, Minsky and Papert focused their analysis on a single line of neural network research. Frank Rosenblatt's seminal book on perceptrons. From this analysis, they drew and actively promulgated sweeping conclusions about the entire field of neural network research, indeed about how everyone should attempt to theoretically analyze biological intelligence. It is well known that these conclusions did not favor neural network research. Everyone who managed to work on neural networks in the 1960s and 1970s can attest to the dampening effects of Minsky and Papert's anti-neural network ardor.

Even in 1969, however, *Perceptrons* represented only one line of research in the neural network approach to understanding biological intelligence. This fact takes on unexpected significance in evaluating the opinions that Minsky and Papert express in their expanded edition.

In the expanded edition, Minsky and Papert again focus their analysis on a single line of neural network research—the popular PDP books edited by McClelland and Rumelhart. As in 1969, Minsky and Papert again draw sweeping conclusions about the entire field of neural network research from their discussion of a single line of work. Again, many of their conclusions do not generalize.

In addition, there is now an ironic historical twist. Minsky and Papert recommend future directions for the development of neural network architectures that they consider especially promising. They call this type of architecture the society of mind. The

irony lies in the fact that the neural network research that they swept away with in *Perceptrons* in the 1960s was already starting to develop, in a technically rigorous fashion, such an architecture. Moreover, many of the technical breakthroughs in neural network research that have occurred during the past two decades were breakthroughs in the design of component modules for a society of mind architecture. I know this fact from my personal experiences because I have participated in making many of these technical discoveries over the past two decades, and I witnessed the intellectual indifference and political hostility of Minsky and Papert to these discoveries when I was a professor at MIT from 1967 to 1975. Even today, Minsky and Papert still seem unaware of these developments.

In fact, Minsky's recent book on the society of mind can be viewed as an elementary nontechnical introduction, using his own terminology, to some of the concepts that had previously been mechanistically developed with technical rigor and much greater depth within the neural network literature. Thus, in predicting the future in their advocacy of society of mind architectures, Minsky and Papert have revealed their unfamiliarity with the past. Their advocacy unwittingly admits that they have made a serious error in their assessment of past neural network research, even as they proceed to make equally sweeping—but on balance more favorable—assessments of current neural network research.

The first page of the prologue is studded with sweeping statements that are not supported by the facts. "Little of significance has changed since 1969. . . . There has been little clear-cut change in the conceptual basis of the field. . . . The spirit of connectionism seems itself to go some-

what against the grain of analytic rigor." These sentences illustrate the authors' willingness to express strong conclusions about topics with which they are unfamiliar and, thereby, weaken the authors' credibility as analysts of neural network research

Minsky and Papert also portray a "war between antagonistic tendencies, called symbolist and connectionist." In reality, among the types of specialized neural architectures that Minsky and Papert would include in their society of mind, neural network examples already exist that carry out the types of symbolic, logical, serial, discrete, localized, and hierarchical operations which Minsky and Papert have attributed to the symbolic (that is, nonconnectionist) movement. In fact, a key research theme of neural network research is to explain how symbolic behaviors emerge within a suitably designed self-organizing neural architecture. Examples of such architectures include the adaptive resonance theory architectures that carry out self-organizing pattern recognition and hypothesis testing in response to noisy and nonstationary environments. Gail Carpenter, Michael Cohen, and I have introduced several of these architectures, and many other neural network researchers are further developing and applying them.

In their critique of PDP research, Minsky and Papert have summarized some properties of the back-propagation algorithm that are well known within the neural network field: Its learning is slow, it does not deal well with certain types of noise, and its computational costs do not always scale well. Much more probing analyses and critiques of back propagation have been published within the neural network literature, including analyses of learning instabilities arising from capacity catastrophies in response to too many input patterns or oscillation catastrophies in response to nonstationary input statistics. It has also been realized that back propagation cannot credibly be used as a neurophysiological model because it includes a nonlocal transport of associative learning weights.

However, the thousands of people who are using back propagation are not all misguided or naive. Rather,

they have found many technological applications exist where the number, statistics, and noise of input and teaching patterns can be adequately controlled, and learning can be slowly carried out in an offline setting. Under such conditions, back propagation can learn associative maps that are proving to be useful.

Do the limitations of back propagation imply that all neural networks fail in real-time applications where the number, statistics, and noise of input patterns cannot be controlled, and learning must take place quickly in an online setting? The answer is certainly not. Indeed, the adaptive resonance theory architectures provide examples of this latter type of learning. Moreover, these capabilities are not merely conjectural. Contrary to Minsky and Papert's claims about connectionists' aversion to analytic rigor, such properties have, for example, been proved mathematically by Gail Carpenter and myself for an architecture, called ART 1, that was described in *Computer Vision, Graphics, and Image Processing* (Carpenter, G. A. and Grossberg, S. 1987. "A Massively Parallel Architecture for a Self-Organizing Neural Pattern Recognition Machine." *Computer Vision, Graphics, and Image Processing* 37: 54-115).

Adaptive resonance theory architectures have also provided examples of systems that autonomously solve a problem posed for society of mind architectures by Minsky and Papert. "When should new layers of control be introduced? If managers are empowered too soon, . . . they would all be overwhelmed by infantile ideas. But if the managers arrive too late, they will retard all further growth" (p. 272). This concern is the special case of a design problem that I have called the stability-plasticity dilemma, which has been analyzed in the neural network literature since adaptive resonance theory was introduced in 1976. (See Carpenter, G. A. and Grossberg, S. 1988. "The ART of Adaptive Pattern Recognition by a Self-Organizing Neural Network," in *Computer*, 21(3): 77-88, for a recent discussion of this problem.) Adaptive resonance theory offers a solution by showing how a neural network can dynamical-

ly self-stabilize its learning. It protects its old learning from being washed away by the flood of new experience, yet maintains its ability to refine its old learning or, where necessary, to generate new internal representations. These self-organized internal representations form an ever-expanding, yet globally self-consistent knowledge representation that continues to develop until the system's full memory capacity is utilized.

Currently, many such rigorously developed examples of specialized architectures exist within the neural network literature. Some of these examples are already mature enough to be entering practical applications. Many are providing a foundation for further basic research.

Another historical irony is that several of these society of mind architectures were published by MIT Press in *Neural Networks and Natural Intelligence* at the same time it published Minsky and Papert's expanded edition, which asserts that such architectures have not yet been developed.

What of the future? As in the conclusion to the recent article in *Science* (Waldrop, M. M. 1988. "Soar: A Unified Theory of Cognition?" *Science* 241: 296-298) about Allen Newell's AI computer program, Soar, I believe it is appropriate to quote a great pioneer of neural network research, Warren McCulloch. "Don't bite my finger—look where I'm pointing!" Minsky and Papert have pointed to a future that is already part of the neural network past. The real future of our field will be determined by a technical literature that is already rich in fundamental discoveries and a community of neural network researchers which includes a stunning variety of talents from many disciplines. How our civilization decides to support and develop these resources will determine how completely its technologies will assimilate the benefits of a biological style of intelligence.

Stephen Grossberg is a Wang professor of cognitive and neural systems, the director of the Center for Adaptive Systems, and the director of the Cognitive and Neural Systems Program at Boston University, Boston, Massachusetts 02215

The Development of an Artificial Intelligence System for Inventory Management Using Multiple Experts

Erwin M. Saniga

The purpose of the research behind *The Development of an Artificial Intelligence System for Inventory Management Using Multiple Experts* (Dayton, Ohio: Council of Logistics Management, 1987, 413 pp., \$50.00) by Mary Kathryn Allen was "to investigate the potential application of expert systems to the management of inventory." Allen was interested in determining whether inventory managers use heuristics, whether these heuristics can be modeled using an expert system, and whether the resulting expert system would yield efficient and effective solutions.

As Allen notes, inventory is big business, and she cites statistics attesting to this fact; for example, the average annual investment in business inventories is about 18 percent of the gross national product. Thus, although academics have been interested in the inventory problem for years, it generates much more than academic interest. Literally thousands of articles appear in journals of mathematics, statistics, operations research, and so on, that address the many problems associated with the management of inventory. Most of these articles have to do with mathematical approaches to subsets of the overall problem. We make assumptions, build models, and optimize the models. We then understand the theory about the problem subsets that are structured.

Nevertheless, the problem of inventory management is not a completely structured one, which is why we see the large gap between the theory and what is actually practiced. Inventory managers use heuristics because of the unstructured nature of the problem and because many of these heuristics are extremely robust. In this context, robustness means that this method works better than any other single method over all products in a firm's product line. The perfect example is what the Japanese have labeled just in time (JIT) inventory management.

Thus, it is not surprising that the answer to Allen's first hypothesis—do inventory managers use heuristics—is yes. The problem she chose to investigate is the expert validation of a system's determination of replenishment spares requirements for repairable items for the United States Air Force. Many managers were making this decision, and Allen used the Delphi method to choose a final group of experts whose knowledge would be elicited in constructing the expert system. The nominal group technique was used in the construction of the knowledge base. Allen's approach in these areas of expert selection and the resolution of conflict with multiple experts was thorough, theoretically sound, and practical. These chapters should provide valuable insights for those involved in the development of expert systems.

Allen modeled the heuristics with a backward-chaining shell, M.I., and the final system contained 441 rules. The details of the system are not discussed in any great length, which is disappointing.

Did the system improve the effectiveness and efficiency of decision making? Allen tested this hypothesis in a 2 x 2 x 3 factorial design having two levels of systems (manual and expert), two levels of problem difficulty (simple and complex), and three levels of user experience (novice, journeyman, and expert). The results of the analysis of variance showed that the expert system led to performance improvements of 15.1 to 17.73 percent for complex cases and 7.6 to 10.4 percent overall. The results for solution efficiency were mixed.

Overall, the Allen book is a valuable addition to any expert system developer's library, especially those sections on the selection of an expert and the resolution of conflict between multiple experts. But the problem of inventory management remains. On the one hand, the structured aspects of the problem have been studied in great detail for many years. On the other hand, Allen's work is one of several that address the unstructured aspects of the problem. What is really needed is a more robust model of the problem of inventory management, which means using the broader mod-

els of the academics along with the experience of the practitioners. Because the problem has both structured and unstructured aspects, a hybrid system using models and rules is the approach to take. At least one firm, Computer Logic, Inc. of Wilmington, Delaware, which is strong in both AI and structured models, has developed an operating hybrid prototype along these lines. Others might also be under development.

Erwin M. Saniga is a professor of Business Administration at the University of Delaware, Newark, Delaware 19716.

Computer Experience and Cognitive Development

Mallory Selfridge

Although considerable research bears on the question of how children develop into adults, almost all is substantially removed from the reality and complexity of actual child development. Occasionally, one becomes aware of child development as a subject of study and is encouraged to see it as an important area for AI research and to believe it will one day be explained in terms satisfying to AI researchers. For example, much of Piaget's work has this flavor, and Minsky and Papert's classic 1972 MIT AI Lab Progress Report presented compelling arguments that AI provides a powerful set of techniques for studying cognitive development. In these days of applied AI, I find *Computer Experience and Cognitive Development* (Ellis Horwood Limited, Chichester, England, 1985, 275 pp., \$28.95) by Robert Lawler to be a refreshing and thought-provoking reminder that explaining how children develop into adults is exceedingly important; exceedingly hard; exceedingly interesting; and, possibly, answerable.

Lawler's intent is to present a theory of the "development of mind" within a child. Although "mind" is a broad topic, and development is complex, he succeeds remarkably well. His theories are based on three sources. First, Lawler performed extensive and participatory naturalistic observations of his daughter Miriam as she learned a number of different tasks. Second, he adopts an

approach to AI that draws heavily on Minsky's society of minds and Papert's LOGO laboratory research. Third, Lawler believes that the body of Piaget's work represents the form and flavor which an integrated theory of the development of mind should take and that Piaget has identified many of the underlying processes of development. Thus, Lawler explains his observations of his daughter's cognitive development in computational terms derived from a synthesis of Piagetian theory and modern AI.

Computer Experience and Cognitive Development is, however, a difficult book to read. Although Lawler's descriptions of Miriam's learning are excellent, his descriptions of his theories are at times difficult to follow. They are expressed only in the most general computational terms. I found myself drawing on years of programming experience to understand him. He frequently and distractingly intermixes data, theory, and philosophy. His statements of the goals of various parts of the book and his conclusions are couched in Piagetesque and individualistic terms and often are only vaguely defined. His diagrams require more study than should be necessary, and several confusing typographic errors appear here and there. Given the choice between a sacrifice of clarity and one of content, Lawler's choice of content was correct, but one wishes for both. I continue to think about this book, considering how the author's theories might be expanded, computer models of the data he presents might be constructed, and his ideas might be applied to my own book; this recommendation is the best of all.

Lawler's book draws on something called "the intimate study," in which Lawler spent six months participating in, and observing the details of, his six-year-old daughter's learning in a number of different areas, with less detailed followup observations. For example, Lawler describes Miriam learning to add, debug LOGO programs, and play tic-tac-toe. Excerpts from his observations and the primary data he seeks to explain are presented throughout the book. These excerpts are naturalistic observations of events, learning situations, and teach-

ing situations in which Lawler himself actively participated. Some might dislike Lawler's departure from methodological rigor in this regard, but I believe that Lawler has skillfully and sympathetically captured the essential reality of this development and that no other technique would do as well.

The role of computer experience ascribed by Lawler to Miriam's cognitive development and his own theories is ambiguous, the title notwithstanding. Some have argued that giving children the correct sort of computer experience can act as a developmental accelerator and guidance system, enabling the child to develop more rapidly and with a greater facility in certain modes of thought; superficially, Lawler appears to fall within this camp. Miriam was heavily involved with LOGO programming during the intimate study, and about half of Lawler's observations concern her computer experience. However, the title *Computer Experience and Cognitive Development* might be somewhat misleading; this book is not trying to be a version of Papert's argument that "computers give children access to powerful ideas." Instead, the computer serves two purposes for Lawler. First, records of Miriam's growing mastery of LOGO serve as a kind of audit trail of her cognitive development and provide Lawler with primary data in a number of important areas. Second, Lawler's primary metaphor for cognitive development is, essentially, computer programming: Cognitive development is the programming of the mental computer. In this regard, Lawler lies directly within mainstream AI.

Lawler explains Miriam's development in Piagetian terms such as "the equilibration of cognitive structures," "genetic epistemology" (which seeks the sources of knowledge in prior knowledge), and "the articulation of complementary roles." Unlike Piaget, however, Lawler's theories give computational meaning to such terms. Lawler argues, for example, that Miriam learns to add by first learning specific procedures which can be applied only to certain objects: She learns to add coin values in a particular way, for example, and angle and

range values in the context of the cue in the LOGO billiards game in some other way, using various cognitive predecessors. General counting expertise develops when several different procedures are executed and produce different answers for the same problem. This unexpected difference of opinion by two procedures within what Lawler views as a small society of minds is what catalyzes development. When this occurs, the two procedures are embedded within a more general control structure that decomposes the problem into parts, selectively invokes each procedure for the appropriate part, and combines their results to solve the problem as a whole. This process is explicitly equivalent to a child writing small LOGO procedures to, say, draw a picture. Lawler explains Miriam's learning to add in terms of genetic epistemology but supplies an explicit computational foundation for this notion, which is lacking in Piaget's work.

The acquisition by Miriam of the ultimately correct procedure for adding does not develop, as might be expected, as a result of a generalizing of the various context-specific addition procedures or a complex "case statement" indexing the context-specific procedures for every possible context. Rather, Lawler argues, learning the correct addition procedure is not unlike the original learning of the context-specific procedures as the combination of cognitive predecessors and, indeed, requires significant amounts of straightforward instruction. However, without the prior development of the context-sensitive addition procedures and their assembly into higher-level structures, Lawler adds, the child would not be able to understand that this new addition procedure she was learning was a generalization of the context-specific procedures she had learned. Moreover, she would not be able to understand the abstract representations within the general procedure and the abstract purposes of the steps in the procedure and would not be able to learn the generalization. Thus, Lawler argues that Miriam's earlier learning provided a kind of cognitive scaffolding without which more advanced learning could not take place.

Lawler's descriptions of child learning and his use of a computational metaphor to describe this learning are extensive in scope, and he presents data on a number of different areas. First, he describes how his son Robby learns to draw pictures using the LOGO drawing program. He proposes that the use of prior knowledge of tinker toys (the tinker-toy "microview") enabled Robby to understand how to command the LOGO turtle to correctly make the triangular roof of a house and, thus, contribute to the development of the LOGO drawing program microview. He explores how Robby's context-specific knowledge in one area is required in order to understand and acquire knowledge about another area, how problem solving is governed to some degree by fortuitous matching of the results of an action and an inactive goal, and how the child applies old procedures in new situations by making constants variables in his procedures.

Lawler then addresses a very Piagetian idea: the invariance of stage development. According to Piaget and most of psychology for that matter, children's cognitive abilities progress through a series of stages, and the order of these stages is more or less invariant. Lawler explores this concept through the use of LOGO programming and, particularly, Miriam's learning to debug her own LOGO programs. Lawler argues that certain elements of advanced, formal thought can exist in certain areas, such as LOGO programming, prior to less advanced concrete thought; LOGO programming can be instrumental in the emergence of formal thought; and formal thought developed within the LOGO context can be generally applied to other areas of the child's life. Lawler concludes that the stage model paradigm requires substantial revision, and "that formal thought, in the Piagetian sense, is a competitor with concrete thought, not an emergent from its perfection" (p. 23). In addition to its primary focus on the issues of stages, Lawler also considers the role of the computer as developmental accelerator and guidance system and discusses the child's acquisition of the debugging metaphor while learning to program and use of the

metaphor in learning in a variety of areas. He concludes that learning LOGO programming (and debugging in particular) is important for Miriam, a conclusion that seems quite plausible. Unfortunately, however, Lawler does not discuss the extent to which Miriam's development could be accelerated and guided by the cooperative and concentrated attention of highly intelligent adults who want to introduce her to exciting, advanced technology primarily designed for fun.

Frequently, there is more to what Lawler presents than what he discusses; the best example of this lies in Lawler's discussion of how Miriam learned to play tic-tac-toe. Now, I had always thought of tic-tac-toe as a pretty simple game that makes a good low-level programming assignment. One can take a Samuels-type approach to learning to play tic-tac-toe or a concept learning approach or a learn-a-new-rule-when-you-lose approach, but I never thought you could learn much from studying tic-tac-toe. After reading Lawler, I have decided I am quite wrong. His description of how Miriam learned to play tic-tac-toe reveals an extremely complex process. To me, the description of Miriam learning tic-tac-toe was a forceful and vivid statement of how complicated children's learning really is and how irrelevant abstract simplifications of the learning process can be. Although Lawler does no more than present his observations and a high-level theory to explain them, it is clear that a computer model of how Miriam learned to play tic-tac-toe could be developed; this discovery is the most exciting of all.

What is important to Lawler in this learning process is how (1) learning to play tic-tac-toe revealed the importance of an "other" in development, (2) how it revealed the importance of the child's ability to think from the point of view of this other, and (3) how the development of the ability to manipulate her own mental environment and think about her own thoughts is central to overall development. Without social interaction, child development does not proceed, and Lawler's descriptions of this interaction in a specific case offer a framework to guide the investigation of

other cases. Lawler also uses this learning example to investigate the development of hierarchies of microviews from a different perspective. He discusses learning to play tic-tac-toe as a process very similar to learning to add: Isolated microviews of context-specific knowledge are coordinated and modified to encompass more and more of the game. As with Miriam, a central factor in this coordination and modification is the ability to mentally play against oneself, and to explore and learn strategic play by restricting the mental opponent simply to tactical play. These propositions seem correct, and the identification and description in computational terms of the vital social role and the importance of reflective thought in child development, although not new, was presented in an entirely unique, detailed, and computationally meaningful way.

Lawler concludes with an interesting thought. Throughout the book, he explains the development of cognitive abilities in terms of the combination and modification of earlier abilities in response to various experiences. The question to ask is: What are the original cognitive abilities? Lawler believes the original cognitive abilities are derived from five "sensorimotor subsystems," each organized around five major body parts. That is, Lawler proposes cognitive development has a specific biological origin. This proposal is interesting and provides a foundation for the later development discussed in the rest of the book.

Overall, Lawler's *Computer Experience and Cognitive Development* provides a coherent view of child development and learning in terms familiar to AI researchers. One can dimly discern the outlines of computer programs that would model Miriam's learning within Lawler's descriptions and theories and herein lies the central contribution of Lawler's book. You read it and start to think that he has been observing and describing some fundamental aspects of the development of intelligence and, furthermore, that a computer model of this development might be within reach. I think *Computer Experience and Cognitive Development* could be

an important book for AI, notwithstanding a certain degree of stylistic opacity, because it represents a high point of AI research into child development and suggests a methodology and a theoretical approach to computer models of child development that capture its essence. I think the development of computer models embodying Lawler's theories would be significant research, and I hope that Lawler or one of his students will develop such models. A computer program that learns to play tic-tac-toe the way Miriam did or learns to add the way Miriam did would really be an achievement.

Mallory Selfridge is an associate professor in the Department of Computer Science and Engineering at the University of Connecticut, Storrs, Connecticut 06268

The Rise of the Expert Company

Robert A. Chalmers

The authors of *The Rise of the Expert Company* (Times Books, New York, 1988, 322 pp., \$19.95), Edward Feigenbaum, Pamela McCorduck, and H. Penny Nii, have given us an absorbing collection of tales about the successful integration of expert systems into mainstream industry. They herald these success stories in terms every manager understands: savers of time and money, improvers of quality and performance, and preservers of the knowledge of key people. It shows these systems are here now, they are real, and they work. It also makes expert systems understandable, no longer a toy from the laboratory but a technology already at work. This approach makes this book ideal for introducing expert systems to people in a position to put them to work.

At first glance, this book does not offer much to attract the AI community. All the expert systems discussed have probably already been described at AI gatherings, some perhaps many times. Thus, it would be surprising if the AI world was really excited by another nontechnical book on expert systems, even with Edward Feigenbaum as one of its authors. The book's natural audience is clearly the managers of organizational units doing

any form of knowledge work.

No, the value of *The Rise of the Expert Company* to AI people is not in their own enlightenment but in their use of it as a powerful sales tool for introducing the technology to potential customers. Tom Peters, coauthor of *In Search of Excellence*, states in his foreword: "I conclude that any senior manager in any business of almost any size who isn't at least learning about AI, and sticking a tentative toe or two into AI's waters, is simply out of step, dangerously so." The middle and upper managers of the world, then, are people who should read this book. However, the people it is really for, the people who should be buying it in quantity to press into the hands of these managers, are the would-be champions and the would-be builders of expert systems.

The persuasiveness comes from good design and good writing. The brief but potent foreword is followed by a main text that consists mostly of engaging narratives about how some twenty-odd companies came to possess expert systems. A lot of first-person insight is given into how these systems were originally sold and developed, with emphasis on the resulting benefits. Scattered between these fun parts is enough supporting background information to make it all understandable. A chapter is devoted to knowledge engineering, distinguishing between it and software engineering. Another chapter is devoted to contrasting the ways different organizations have found to phase in the use of expert systems, and the variety is surprising. The overall perspective is not that of the history or principles of AI, it's one of the productivity of knowledge workers, of global competition, and of economics and profits.

The style of the writing is direct and engaging and benefits from the generous use of quotation. There is also much of the flavor of McCorduck's excellent earlier work, *Machines Who Think*.

The book concludes with an appendix by Paul Harmon, who provides data on 140 expert systems in use, including all those discussed in the body of the book. For each, he identifies the customer and the devel-

oper, what it does, the tools used, a contact name and phone number, and the system's availability—about one-third of them are for sale. Although not a large enough sample to draw statistically meaningful conclusions, the material helps convey the diversity of working program domains as well as of the organizations employing them.

I think the book could have been improved in some ways. I regret the absence of an index in light of all the factual material it contains. I felt the least satisfied with the ending of the main text. It probably wasn't easy deciding how to end this book. We are given a summary of the gains expert systems are capable of producing, the technical challenge of the expert system of the future, and some other final thoughts. However, the big pluses are diluted by some serious problems introduced only in these final pages. It is important material, but it belongs earlier in the book. I feel the lack of a strong positive closing, the salesman's clincher, is a definite weakness in this otherwise highly effective presentation.

Also, any manager given a good persuasive pitch for using money wants to hear the risks presented with the benefits, even if not with the same fervor. However, little is said about failed attempts among the successful case studies (Du Pont's Ed Mahler is quoted as claiming [only!] a 90 percent success rate for over 200 small systems, and at least one system, we are told, had more than one start.)

Finally, only a paragraph is given to a serious potential concern—the displacement of workers. One cannot say a lot in the abstract, but after all, this book is a collection of real case histories. The occurrence or absence of job displacement, or the steps taken to forestall the problem, could have been included in each case.

These are but small flaws in a large asset. One more, however, has just occurred to me: The last copy I loaned out has been gone much too long. I can't keep this book on my shelf.

Robert A. Chalmers is a staff scientist in the Applied Physics Laboratory at Lockheed Research Laboratory, 3251 Hanover Street, Palo Alto, California 94304