Book Reviews

A Step toward the Understanding of Information Processes

Perceptrons. An Introduction to Computational Geometry. MARVIN MINSKY and SEYMOUR PAPERT. M.I.T. Press, Cambridge, Mass., 1969. vi + 258 pp., illus. Cloth, $12; paper, $4.95.

This is a great book. To understand this judgment, and why I am willing to make it at so early a date, is not so simple. For the book is many things, all of which must be exposed before the grounds of judgment are clear.

Let me first describe most narrowly the content of the book. It is a contribution to the theory of perceptrons. Mathematically considered, a perceptron is a predicate that can be represented as whether a weighted sum of other predicates exceeds a specified threshold. That is, we are given a space R (usually called the retina) within which we can distinguish subsets of points, X (usually called pattern instances). We wish to consider predicates, $\psi(X)$, which are true if X is a member of a family of classes (the pattern) and false otherwise. For instance, $\psi(X)$ should be true if and only if X is a circle. Or, $\psi(X)$ should be true if and only if X is a connected figure. Let there be a set of predicates, call them ϕ, which one can think of as elementary measurements on the space R. Then a perceptron is a predicate which can be represented in the form:

$$\psi(X) \text{ is true if } \sum_{\text{all } \phi} \alpha(\phi)\phi(X) > \theta$$

$$\psi(X) \text{ is false if } \sum_{\text{all } \phi} \alpha(\phi)\phi(X) \leq \theta$$

where the coefficients, α, and the threshold, θ, are real numbers and the values of the ϕ predicates are taken to be 1 if $\phi(X)$ is true and 0 if it is false.

Perceptrons represent a particularly simple class of pattern recognizers (or equivalently, decision elements). They collect evidence from a number of separate sources, the $\phi(X)$, weight it in some fixed way, the α, and ask if the evidence adds up to enough, θ, to warrant saying that X is an instance of the pattern (equivalently, deciding yes). Although this corresponds to the oft-expressed intuitive notion that judgments are made by "weighing the evidence," it must be made clear that perceptrons are an extremely restricted class of decision devices. In most real decisions there is much exploring of consequences, returning for new information, redefinition of the situation, and so on. None of these processes find expression in the perceptron, as formulated. Nevertheless, perceptrons still constitute a nontrivial type of decision element, and—as Minsky and Papert note—if we cannot understand the behavior of perceptrons we have little chance with the more complex decision processes.

The book states and proves a large number of theorems about perceptrons. For any interesting theory, one must restrict the elementary measurements (the ϕ), since otherwise the whole burden of the decision could be put on them, the combinational aspect that is the essence of the definition thus being bypassed entirely. Two restrictions are proposed: *diameter-limited* perceptrons, in which the points on which a ϕ depends must all lie within a circle of given diameter (though the whole collection of ϕ can cover R many times over); and *order-limited* perceptrons, in which the number of points on which a ϕ depends must be less than a given number (though the points can be located anywhere on the retina). Both restrictions fit an intuitive notion that the ϕ are somehow simple, limited and local predicates, so that the act of combining evidence is really left to the summation.

To each way of restricting the ϕ there corresponds a separate theory, and both theories are developed in the book. The order-limited theory is, so far, the more interesting. Typical ultimate results are that there is no perceptron that can recognize when a figure is connected, as opposed to being disconnected. This holds for both diameter-limited and order-limited perceptrons, though the proof for the first is direct and for the latter quite complex. In general the results are of this negative character. For instance, it is possible for there to be perceptrons of order 1 for two predicates, yet no perceptron of finite order that will recognize the disjunction (or, similarly, the conjunction) of the two predicates. In the development of the theory some powerful tools are constructed. Perhaps the most central is the group-invariance theorem, which states that if a perceptron is to be invariant over a (finite) group of transformations on the retina, then there must exist a particularly simple form of the weighted sum (namely, where all coefficients of those ϕ which are equivalent under the group are the same). The power of this theorem arises from the close connection between notions of what is interesting geometrically and properties that are invariant under groups of transformation. Thus the theorem reflects something of the geometry of the retina in the algebraic structure of the perceptron.

Still other results concern the fact that though order-limited perceptrons exist for some classes of patterns, their coefficients (more precisely, the ratio between the smallest and largest coefficient) may be exceedingly large—so large, indeed, that one might as well store the instances directly, since that would require fewer bits than storing the coefficients. There is a chapter on learning in perceptrons in which one considers the ϕ fixed and asks what procedures might discover appropriate weights to do a particular pattern-recognition task. The information from which the weights are inferred is a sequence of instances of the patterns. There is a perceptron convergence theorem which states that a particularly simple form of feedback modification of the weights under the impact of the sequence will indeed find a workable set of weights if such exists. Finally, there is a comparison of the perceptron with various highly serial algorithms for recognizing some of the same kinds of patterns—exploring the time-memory trade-offs between the two.

So much for the direct contents. This is a book on the mathematical theory of a class of abstract devices, and without further enlightenment you might never suspect that it merits the adjective great. For it will be an "insider's" great book and will be read only by a relatively small number of computer scientists and cyberneticists. Unlike a number of other books with which it might be compared—Shannon's little volume on information theory, Von Neumann and Morgenstern's treatise on game theory, or Hebb's *Organization of Behavior*, its intrinsic subject matter does not have the overwhelming breadth to catch the world's eye. Still, the judgment stands.

The work on perceptrons had its genesis, as so much of cybernetics has had, in the work of Pitts and McCulloch in the 1940's in putting forth an abstract model of a neuron, so that one could envision modeling the activity of the brain by networks of logical neurons. In 1958 it received definitive shape when Frank Rosenblatt selected for attention that subclass of neural nets that consists of a retina plus a single set of association neurons which take input signals from the retina, summate and threshold them, and deliver a pulse to response neurons that represent the final decision. These devices Rosenblatt called perceptrons. Central to this work was the idea of training the perceptron by reinforcement procedures while exposing it to a sequence of pattern instances. The subject became both popular and controversial, since these devices were viewed by some as having remarkable powers of self-organization and as being the first true toehold into the development of really intelligence devices.

In the decade of work that followed Rosenblatt's work there occurred essentially three developments in understanding. First, it was realized that the underlying machine was that of linear threshold devices and that these devices had an independent engineering interest. Second, if one considered the measurements of the ϕ as simply a point in n-dimensional space, then the theory could largely be assimilated to that of linear discriminant analysis in statistics. Third, the perceptron convergence theorems provided some understanding of the learning capabilities of the system. However, nowhere in this entire body of work did anything exist on the actual capabilities of perceptrons. What geometrical predicates could they recognize and what ones couldn't they recognize? This was reflected in the complete lack of any results on the *rate* of convergence. Yet these questions are precisely the ones at issue in the controversy about whether in perceptrons one had a major avenue for achieving intelligence machines.

The present book, then, records a major theoretical tour de force in creating a full-bodied theory, almost *de novo,* that tackles these central questions successfully—though of course it does not answer all of them. Except for the chapter on the perceptron convergence theorem, the formulations, questions-to-be-posed, theorems, and proof techniques are original. Those parts that are not original with just the two authors appear to be primarily the work of students and scientific friends who have reacted in a positive way to the theory as Minsky and Papert were developing it. And when this theory came it had almost no intersection with formulations of the past—that is, with statistics and the existing theory of linear inequalities.

With respect to the history, Minsky and Papert have shown the old controversy (is the perceptron a good thing or isn't it?) to be irrelevant. One might have expected as much. For the real point is that such elementary decision processes must be understood, and as deeply as we know how. By their reformulation and by their willingness to acknowledge the past efforts (for example, by entitling their book *Perceptrons*) Minsky and Papert have resolved that particular controversy in a scientifically most appropriate way.

Now consider the larger theme of what computer science is and how it shall be executed. The present book is one of a very few entries into that collection of works which are entirely and exclusively studies in computer science. (Another is the seven-volume *Art of Programming* by Donald Knuth, now with two volumes extant.) The book says, by demonstration, that we in computer science are to understand classes of information processes—of algorithms, if you will. It defines such a class and sets standards for the nature of that understanding. And the class of processes that it considers (perceptrons) is one that comes from somewhere near the center of computer science—one that is not related to reactor calculations, analyses of variance, or crystallographic calculations. Pattern recognition is a phenomenon that has arisen from around the computer (more generally, information-processing technology), and it is part of our proper subject matter.

But there has always been a theoretical side to computer science. Indeed, it existed as the well-defined field of automata theory (and recursive function theory) long before computer science came into formal existence. Is not the work under review simply one more effort in that genre? The answer is both no and yes. There has always existed a most uncomfortable gap between automata theory and the heartland of computer science—programming and particular algorithms for accomplishing specific tasks. There is an easy explanation for this, namely, that automata theory and recursive function theory have been shaped by a concern with ultimate computability—with whether any process could be found that would compute certain broad classes of functions. Many calls have gone out for a theory that deals with the expanse and complexity of actual algorithms of practical interest. But such a call is easier given than heeded.

The present book is a very long step in the right direction. But I've said it wrong. The present book starts from the right set of problems—to understand a very specific class of algorithms in the way computer scientists want to understand them—and then takes as small a step as possible away from that central goal in order to get interesting results. So, although the results in the book are often "theoretical," in the quaint sense of the term that opposes to "practical," the theory as a whole has the right shape and the right concerns. In this sense the answer is no, the book is not of the genre of automata theory. It is worth remarking here that Minsky and Papert are not mathematicians recently arrived on the scene just in time to formalize it appropriately. They are responsible for the Artificial Intelligence Project at M.I.T., which has been one of a few centers of empirical work in the heuristic programming of computers over the last decade.

However, parts of automata theory are also moving in this direction, in particular that subpart called complexity theory, which tries to characterize how complex various classes of computations are. Thus the gap in general is not so large as it once was. In this sense the answer is yes, the book is not so far from what some automata theory is becoming.

We are not finished with fundamental issues of concern to computer science that cluster around this book. Another is the distinction between performance and learning. There is throughout that part of computer science that deals with what is known as artificial intelligence a fascination with learning. For instance, the first question that is likely to be asked about a chess program, no matter how good, is "Does it learn?" This feeling that learning is somehow the touchstone of intelligence is held not just by novices in the field but by many seasoned workers as well. (It may have something to do with the fact that if a machine learns, then it will not have been the case that it was instructed by someone, as appears to be true of a programmed digital computer, and thus is not open to the charge that it will do only what its programmer has told it to do.) The history of the perceptron itself provides an example of the effects of such a view, where through ten years of intensive investigation (several hundred published papers, at least) almost nothing was learned about performance and almost all attention was devoted to learning. The present book carries the strong message that one must understand the performance characteristics of a class of systems before the learning issues can take on meaning. For learning is always the selection of a particular performance program within a space of possible performance programs. Understanding learning requires, first and foremost, understanding the structure of this space. Again, the book carries this message, not so much by saying it (though it does that, too, here and there) as by attending to the theory of the performance of perceptions and showing how fruitful that theory can be.

A final consideration is the style of the book. It is essentially mathematical, the mathematics in it is not superficial, and the book has the standard backbone of a small set of definitions and a large series of theorems and proofs. But beyond this point it parts company with most mathematical treatises. For the authors are much concerned with the heuristics that guided them, with the clues in the environment that led them on, with the barriers that held them up, with the nature of the reasoning that has been transformed into proof method, and with possible conjectures and concerns about the future shape of the theory. The book continuously deals with these matters, and one comes away with a feeling of having taken a guided tour. It is a fine book for individual study.

Perhaps, now, the grounds for my assessment are clear. It would do no good to repeat them. All that I have said is favorable. The book of course has its quota of minor flaws, but I have no urge to temper my judgment on the larger issues surrounding the book with a few niggles in an attempt at judiciousness.

I should remark, perhaps, that I am not an unbiased witness, although I trust I have kept my wits about me in examining the book. For I share with Minsky and Papert a common view of the appropriate shaping of computer science into a disciplined field of inquiry. And I see no need to give other than my true assessment of the potential role of this book in that shaping.

ALLEN NEWELL
Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213

A Genus of Small Mammals

Biology of Peromyscus (Rodentia). JOHN A. KING, Ed. American Society of Mammalogists, Stillwater, Okla., 1968 (available from Bryan P. Glass, Department of Zoology, Oklahoma State University, Stillwater). xiv + 594 pp., illus. $15. American Society of Mammalogists Special Publication No. 2.

An introduction by W. Frank Blair sets forth clearly the rationale for this fine work. "The diversity of life is so great that we can never hope to answer all of the questions we might pose for all of the kinds of organisms, so we look to detailed studies of representative taxa for elucidation of the principles governing diversification and maintenance of diversity. . . . Ideally, the taxon should have a good fossil record . . . should have a considerable amount . . . of diversification in living forms so that various stages of the evolutionary process are adequately represented," and should be "one that is actively evolving so that the mechanisms of evolution may be investigated as dynamic processes. Ideally the organisms should be obtainable with relative ease and in adequate numbers, and they should be amenable to laboratory culture. The North American genus of cricetine rodents, the white-footed mice or deer mice of the genus *Peromyscus*, about which this book is written, come as near to fulfilling these requirements as one might hope for in a taxon of small mammals."

Many attributes of white-footed mice contribute to their usefulness in investigating the dynamics and behavior of natural populations. They are of fitting size; most of them are easily trapped alive; some live readily in nest boxes; and they are ubiquitous in North America, probably being the most widely distributed native small mammal and ofttimes by far the most common one. Several species thrive in captivity, and it is thus possible to investigate modes of inheritance of genetic traits, physiological attributes, and behavioral characteristics.

The paleontology of *Peromyscus* is treated in this book by Hibbard and the classification by Hooper. Of the impressive volume of information on the genus, the largest share concerns *P. maniculatus* and *P. leucopus*. We may know more about *P. maniculatus* than about any other wild mammal. Dice, one of the fathers of peromyscology, contributes the chapter on speciation, considering exclusively the splitting of a species into daughter species. "Habitable areas for mice are rarely mouseless for long," relates Baker in his section on habits and distribution. White-footed mice are among the first mammals to invade disturbed habitats, even having been found living in the shelter of year-old Volcan Paricutin lava. Klingener tells us that understanding of the anatomy of white-footed mice, though poor and fragmentary in comparison with knowledge of the dog and laboratory rat, is rich and extensive in comparison with that of many other genera of rodents. Perhaps the greatest flaw in our understanding is the inability to correlate most of the observed structural differences in the genus with differences in function and behavior. In an encyclopedic chapter of 106 pages, Layne reviews morphological and behavioral development and growth, including unpublished work of his own. Knowledge of embryology is limited to a single species, *P. polionotus*, whereas data on at least some aspects of postnatal development and growth are available for 12 species.

Whitaker summarizes information on large parasites and points out subjects for further research. White-footed mice "are not to be considered as major sources of human diseases; indeed, they appear to be very clean little animals." Endocrinology is reviewed by Eleftheriou, whose work with *P. bairdi*, together with that of others, represents the first systematic attempt to clarify the role of the amygdala in secretory