

use of probability and the theory of stochastic processes, an example of material that would be helpful in handling real-life problems, is given virtually no attention.

Restricting the number of potential readers is unfortunate because an interdisciplinary view of the world around us must be developed. This book should have been written to show a scientist with a good mathematics background how to do modeling and simulation. Scientific research needs more people trained in system concepts, people trained to understand and apply the Weltanschauung of system theory. Indeed, the recent recommendation for science education that came out of the Science for All Americans study, sponsored by the American Association for the Advancement of Science, emphasized an interdisciplinary approach to scientific concepts. By limiting the technical accessibility of this book, the author has not helped us address the need for training scientists in the use of interdisciplinary tools in scientific research.

The text will be difficult to use in self-study; a great deal is expected of the reader. Explanations of the models are frequently difficult to follow. A lack of answers for the exercises makes assessing what one has learned difficult. For a graduate-level course for well-prepared students learning from a well-trained professor, though, this text might be useful. It has no real competition among books on applications of mathematical modeling. However, in my opinion, there is still a need for a book on mathematical models that is accessible to a wider readership.

*Mark E. Lacy is manager of computational biology at Norwich Eaton Pharmaceuticals, Inc., NY 13815.*

## Sparse Distributed Memory

*Terry Rooker*

*Sparse Distributed Memory* (Cambridge, Mass.: MIT Press, 1988, 155 pages, \$24.95, ISBN 0-262-11132-2) is an interesting little book in which Pentti Kanerva describes a fascinating theory of human memory. Almost as surprising as the elegance of the theory is the length of the presentation: He uses only 120 pages to describe the theory, with the more formal mathematical results given in 25 pages of appendixes. With consid-

eration to the insight required to develop these ideas, the brevity of the description is remarkable. This issue of brevity is important because it makes accessible a number of ideas that many people will find interesting.

Sparse distributed memory (SDM) (I use the initials to distinguish the theory from the book) is an idea that has been developing for some time. Kanerva published his dissertation in 1984. Since then, the basic theory has been used in a variety of ways as a model of human memory and a model for a new style of computer memory. Publication of this book will bring the theory to the attention of a much wider audience. For this audience, there are two important aspects to the theory: It is inspired by the study of brain physiology and is able to explain many of the observed functions and behaviors of the human brain. Unlike other neural models, SDM easily scales to large vector sizes.

At various points in the book, Kanerva describes how he developed his ideas in an attempt to provide a computational description of structures in the brain. This effort seems to have been successful: SDM can be mapped onto physiological structures, a feat that many alternatives cannot duplicate. Many neural models only duplicate a style of computation and are not intended as models of brain functions. Kanerva admits that SDM does not begin to capture the complexity of the brain, but it is at least a closer approximation than many other such efforts.

A problem with many neural models is their computational expense. They work fine with small networks but do not scale well with an increase in the number of nodes. Hardware implementation does not help because there is no obvious way to provide communications paths for a fully connected network. SDM, however, does not perform well with small numbers of nodes. It requires a million nodes to realize the beauty of the system. Kanerva describes how hardware implementation can be achieved with the use of numerous counters. Because each counter location does not need information from other nodes, there is no communications bottleneck, and hardware implementation is much simpler than for other neural models.

SDM uses vectors in  $n$ -space as input. Although the vectors could be

in other bases,  $n$ -bit binary vectors are used in the book. Input are matched to target vectors by simply finding the target closest (in terms of Hamming distance) to the input vector. If  $n$  is sufficiently large (thousands or even millions of bits long), this vector space exhibits some interesting properties. Most important, the bulk of the vector space will be clustered around the midpoint between any two vectors, in much the same way that most of the surface of a sphere is located midway between any two opposite points on the sphere.

To directly implement such a scheme would require  $2^n$  memory locations. For the large  $n$  required, such an implementation is impossible. To avoid this problem, Kanerva proposes using only memory locations that actually have values stored in them. As he points out, a century has fewer than  $2^{32}$  seconds; thus, it seems likely that only a small fraction of vectors in a 1000-bit vector space will actually be stored. Therefore, we can use many fewer locations than the size of the vector space. Such a sparse memory could feasibly be implemented and, depending on the interpretation of the input, used as an associative memory, for best matching, or for any of a variety of other applications.

A separate and interesting part of this book is a critique of perceptron convergence learning. Many connectionist models rely on the perceptron convergence theorem. Kanerva outlines one problem with perceptron convergence to show why he rejected this approach. Briefly, perceptron learning requires an outside reference to discriminate between different classes that are being learned. Thus, the discrimination has already been made, which is counter to most physiological models of neuron function. Many perceptron convergence-based algorithms are used for classification applications for which training data do exist. In these cases, Kanerva's critique does not apply. It does, however, make it clear why he rejects convergence learning as the basis of a physiological model.

SDM's strength as a physiological model rests in part on its simple, elegant theory. Its ability to scale well and the lack of any complex implementation requirements also lend strength to its claim as a model. Brain physiology is notorious for the

You are cordially invited to become a member of the  
AI Community's principal scientific society:  
**The American Association for Artificial Intelligence**

CONSIDER THESE BENEFITS . . .

- **AI Magazine**  
A quarterly publication (called the "journal of record for the AI Community" by High Technology magazine) devoted to the entire field of artificial intelligence.
- **AAAI Membership Directory**  
An annual roster of AAAI members sharing an interest in artificial intelligence.
- **Proceedings of the National Conference on Artificial Intelligence**  
An annual publication consisting of state of the art research papers presented at the AAAI's annual conference. (This publication is included in conference registration; AAAI members not attending the conference may purchase the Proceedings at a reduced rate.)
- **Technical Program of the National Conference on Artificial Intelligence**  
Highlighting examples of significant contributions in the development of state of the art in AI practice. (AAAI members receive a substantially reduced rate on conference registration fees.)
- **Exhibit Program of the National Conference on AI**  
Affording attendees the opportunity to view the latest in AI products, services, and research.
- **Tutorial Series**  
On state-of-the-art topics. (AAAI members receive a reduced rate on tutorial registration fees.)

- **Substantially reduced subscription rate for Artificial Intelligence**
- **Assistance in formation of local and regional AI groups**

To take advantage of all these benefits, complete and send the coupon on the right together with your check (in U.S. funds only) for \$25.00 (\$40.00 foreign) to the address below (or call 415/328-3123 for further details). Send to:

**AAAI Membership**  
445 Burgess Drive • Menlo Park, California 94025

Name \_\_\_\_\_  
Company/Mail Stop \_\_\_\_\_  
Address \_\_\_\_\_  
City \_\_\_\_\_  
State/Zip/Country \_\_\_\_\_

simple computational abilities of neurons, and there are on the order of  $10^{11}$  neurons. Both these facts run counter to other connectionist models but easily fit SDM.

*Sparse Distributed Memory* will be of interest to anyone doing research in neural models or brain physiology. As the theory is refined, the book will also be of interest to those trying to find applications for neural models. Finally, it will be fascinating to anyone who is even slightly curious about human intelligence and how it might arise from the brain.

*Terry Rooker* is a graduate student at the Oregon Graduate Institute. His research interests include computer speech recognition using neural networks.

## Pattern Recognition

*Scott W. Shaw*

*Pattern Recognition* (New York: John Wiley and Sons, 1987, 144 pages, ISBN 0-471-61120-4) by Mike James is a concise survey of the practice of image recognition. The title might be somewhat misleading because the term pattern recognition generally refers to the statistical or syntactic classification of abstract entities. This book, however, discusses techniques

for processing and extracting information only from images. The intended audience has an undergraduate education or less, with little math or computer experience. The goal of this book is to give the average personal computer user some exposure to computer image analysis and its benefits. Toward this end, the book is a success. However, it is not appropriate as a college text or reference work. The text is well composed and reads easily, but the depth of coverage is disappointingly thin. The text might be a good introduction to the subject for a technical manager whose time is limited or a hobbyist with no formal background in the subject.

The book begins with an introduction to images and how they are acquired, digitized, and stored. In this discussion, James points out that images are functions from spatial coordinates to grey levels. He then cautions the reader not to view image analysis as a branch of function theory. This bias against formal mathematics persists throughout the book.

The second chapter provides a brief introduction to pattern recognition (as it is more commonly known), that is, clustering vectors and constructing discriminant functions. This

short outline covers everything from the definition of a probability to the perceptron learning algorithm.

The third chapter deals with identifying local image features, such as lines and edges, by template matching and linear edge operators. The Sobel operator, various gradient approximations, and the Laplacian are presented with short, heuristic descriptions. This discussion is followed by a chapter on frequency domain operations, including an introduction to the Fast Fourier Transform, spatial filtering, and a brief section on Fourier descriptors. The text is augmented throughout with example images illustrating the concepts being described. These images are indispensable for providing the reader with an intuitive feel for the operations. However, simple visual inspection of image-processing results can lead to a false sense of confidence in computational algorithms. An accompanying mathematical analysis is often far more enlightening.

The fifth chapter is devoted to image regions and boundaries. A particularly useful section motivates the use of variable thresholds for region segmentation. This chapter contains