AI Magazine Volume 11 Number 2 (1990) (© AAAI)

Book Reviews

Alternate Realities: Mathematical Models of Nature and Man

Mark E. Lacy

In his new book *Alternate Realities: Mathematical Models of Nature and Man* (New York: John Wiley and Sons, 1989, 493 pages, $34.95), John L. Casti gives us an impressive, up-to-date look at several areas of mathematics that are being applied to the study of biological and sociological systems. These areas, including cellular automata theory, catastrophe theory, nonlinear dynamics and chaos, game theory, and control theory, are finding use on the frontiers of scientific research. Although these areas and their applications are described in various other sources, both on the level of a scientist and a layperson, I know of no other book that brings them all together to show how they can be used in scientific research. However, this book suffers from being written for mathematical specialists and, therefore, limits the potential readership. An opportunity to educate more scientists in the use of mathematical models is regrettably missed.

Each chapter of this well-organized book deals with a different class of models and includes examples to illustrate the author's points. Substantial sections of discussion questions and problems and exercises can be found at the end of each chapter. This is a nice feature that many related books lack. The discussion questions are interesting and thought provoking if somewhat frustrating because they really do require discussion instead of just personal reflection. The problems are typically mathematical (show, prove, and so on) rather than computational (calculate, write, and algorithm or program), and no solutions are provided. My preference would have been for more computational problems to help the reader understand how to directly apply the concepts.

In the final chapter, Casti provides a brief discussion of how the theory of models fits with the broader theory of representations. Unfortunately, the mathematical models covered in this book are not directly tied to AI, where the issue of representation is premier. In the chapter on brains, minds, and mechanisms, the author states that he deliberately "omits any consideration of intelligence," choosing instead to focus on properties of observability and reachability and models based on cognitive and behaviorist schemes. A separate chapter discusses the use of cellular automata models but omits any mention of the use of these models to study systems of excitable cells (for example, neural tissue, cardiac tissue) or of the similarities and differences between cellular automata, interacting particle systems, and neural network models. Neural networks, a topic of much current interest, are not discussed in this book.

Although it would have made the book somewhat longer, it would have been useful to include both a beginning section on general system theory and a liberal sprinkling of discussions on computer simulation. It seems to me that without the former, it is difficult to appreciate the importance of models, and without the latter, it is difficult to apply them. The successful system scientist needs to know not only the mathematical models available but also the philosophy behind the use of models (of all kinds) and how the behavior of a model is studied through simulation. Those models whose behavior can be analytically studied without simulation are few and far between. It would have been helpful as well to include some explanation of the relationships between the different kinds of models and how one goes about choosing an appropriate model for a selected problem. Such an explanation would have improved the cohesiveness of the author's presentation.

Although they are not part of the focus of this book, Casti's remarks on the function of the system scientist are nevertheless notable. This function, he says, is to be the holder of the magic keys, "the abstract encoding/decoding operations" that enable us to connect the mathematical world back to the real world through prediction and connect the real world back to the mathematical world through observation and measurement. The system scientist need not be an expert in mathematics or a particular scientific discipline; the scientist instead bridges the gap between the two.

Thus, I arrive at the major obstacle with this book. I expected the book to be written for the system scientist or a scientist who wishes to use systems methods, such as modeling and simulation. However, it seems to be written for mathematicians instead. In conflict with his statement that the systems scientist need not be an expert in the mathematical world or in the real world, Casti states in his preface that the reader is assumed to have a working (note, not passing) knowledge of ordinary differential equations, linear algebra, matrix theory, and abstract algebra. (I suggest that a background in differential geometry and algebraic topology would be helpful as well.) The book is indeed written for someone with such knowledge, but few readers other than mathematicians will have this training.

It is disappointing that this text makes use of a high level of mathematical detail; it will be a formidable obstacle to its use and will unnecessarily restrict the size of the book's potential readership. The reader must be exceptionally knowledgeable in mathematical terminology and methods and must be prepared to wade through pages of esoteric notation. The value of some of the extensive mathematical developments is not clear, whether it is in support of an understanding of a method or in its application. Casti presents material that is so technically dense that it would be difficult to take it and apply it to real-life (dirty) problems. The 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 Pentii 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 consideration 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 in 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