Books

Machine Learning: Paradigms and Methods, J. Carbonell, Editor, A Bradford Book, The MIT Press, Cambridge, MA, 1990. 394 pp. (ISBN 0-262-53088-0).

Computer power has been estimated to have grown by a factor of one million to ten million in the last forty years, and many believe that the trend will continue. If one stops to think about what computers of the future with ten million times their present power might do, and how they might be programmed, one is drawn to the idea of machine learning—that is, computers using experience to program themselves. This becomes particularly relevant if much of the increased power is achieved by parallel architectures.

That is why this book is particularly appropriate at this time. It is a collection of nine specially commissioned chapters by a total of 18 authors. In the introduction the editor explains that "Each author was asked to write a self-contained article with explicit historical and cross-paradigmatic perspectives for a reader well informed in artificial intelligence, but not necessarily an expert in machine learning."

The book is readable and achieves its purpose. Each chapter offers its own perspective on an aspect of machine learning. A glance through the chapters reveals the degree to which machine learning is a diverse field of endeavor, so the approach to the book was well selected; there is probably no single author who could do justice to this broad topic. While recognizing that work in machine learning is as old as the computer industry itself, the book also reveals how much more there is to accomplish. It is a status report on the state-of-the-art, and an entirely different book will be written ten years from now.

Some of the methods described in the book are being used successfully today. The chapter "Model of Incremental Concept Formation" discusses induction techniques, such as Quinlan's ID3 method for creating decision trees, which has been at the core of several knowledge engineering products. The chapter "Connectionist Learning Procedures" discusses a family of methods, including artificial neural networks, which is being used in financial, marketing, scheduling, and product diagnosis applications.

The chapter "Classifier Systems and Genetic Algorithms" describes some interesting methods that are being actively explored but which are still mostly research topics, such as genetic algorithms. This technique focuses upon the creation of new concepts as well as the evaluation of known ones.

The remaining chapters focus on symbolic learning and representations of knowledge, and report on research results which describe or demonstrate principles but fall short of demonstrating feasibility. To quote from the chapter on "Explanation-Based Learning":

> EBL is still a rather new paradigm and an area of very active research. ... [In] extending the use of EBL a host of new issues arise.

And quoting from the chapter on "Design by Derivational Analogy":

> Design by derivational analogy has turned out to be easier said than done. Clearly much work remains to be done to remedy the deficiencies of current systems and make it a practical reality.

These two quotes could as easily be applied to much of the work reported on in the other chapters, including "Data-Driven Approaches to Empirical Discovery," "A Theory of the Origins of Human Knowledge," and "Creativity and Learning in a Case-Based Explainer."

Certainly, much work needs to be done to determine just what machine learning is. Except for the chapters on connectionist learning and classifier systems, there was scant mention of learning paradigms outside of the purely symbolic arena. Does this mean that there is no work going on in other areas, or that the editor chose not to include it?

As a checkpoint in a changing field, many unanswered questions remain. How does learning relate to training, on the one hand, and creativity, on the other?

Training implies that a structure has been created in the absence of certain important parameters that are determined by the training process. Training is just slightly removed from mathematical optimization and best describes the current results of the connectionist methods.

Creativity, on the other hand, implies leaps of understanding and the identification, not of parameters, but of wholly new structures. Genetic algorithms are moving in this direction although they are not yet there.

Machine Learning spans the vast and largely unexplored distance between these extremes. The reader who picks up this book in the hope of finding answers will be disappointed; the one who picks it up looking for ideas and inspiration will be rewarded.

Richard P. Ten Dyke
Bedford
New York

Note—The books reviewed are those the Editor thinks might be of interest to our readers. The reviews express the opinions of the reviewers.