

A Review of *Machine Learning*

Dennis Kibler

Tom Mitchell states that the goal of his text *Machine Learning* is to present the key algorithms and theory that form the core of machine learning. Not only has Mitchell succeeded in his primary goal, but he has accomplished a number of other important goals.

My major interest in this book is as an introductory text for a first-year graduate course. The goal of the class is to put students on a fast track for doing research in machine learning. After the course is completed, students are expected to have the core knowledge and skills necessary to begin directed research, which entails much more than a knowledge of previous results, such as might be presented in a book that reviews the previous literature. One also needs to learn the research methodology, the research language, and the research goals of the discipline. *Machine Learning* provides these ingredients.

Mitchell uses established terminology to describe previous work and prepares the reader to understand current research. After reading this text, students will be able to understand the literature as well as conference presentations.

The text begins with a clear definition of learning together with significant motivating problems, so that students can appreciate both the value of machine learning and the difficulty of the task. He also puts forth the important questions that machine learning addresses so that students will have a good foundation on which they can build their knowledge. In this way, Mitchell stimulates students to think about important questions, methods for solving them, and approaches for evaluation. Each chapter ends with a set of questions that enrich the under-

standing of the power and limitations of learning algorithms.

Machine learning draws on multiple disciplines. Mitchell provides the necessary background in both statistics and computational learning theory (a chapter on each) so that results from these fields can be understood and applied. He does not go overboard and overwhelm students in these areas. Instead, Mitchell takes the practical point of view. Students are provided with enough information to understand and use results from these ancillary fields.

As was his intention, Mitchell does not attempt to cover all machine learning. Instead, he selects those topics that are most developed or most promising. One unusual quality of this

Machine Learning,
Tom M. Mitchell,
McGraw-Hill, New York,
New York, 1997,
414 pp., ISBN 0-07-
042807-7.

text is that Mitchell always searches for and presents, when available, the underlying supporting theory. The text primarily covers concept learning. Methods using version spaces, decision trees, perceptrons, linear threshold units, artificial neural nets, radial basis function, instance, conditional probabilities, Bayes's networks, rules, first-order logic, and background knowledge are clearly explained. Both experimental and theoretical evaluations are given for existing algorithms,

with indications about where one might expect advances.

Several topics of interest are not covered, including learning in problem solving, combining multiple models, and clustering. Older methods of learning problem solving, such as by acquiring heuristics, macros, or control rules, are not covered, perhaps because this arena has seen less research in recent years. In fact, two active research methodologies for problem solving—genetic algorithms and reinforcement learning—are both introduced by Mitchell. In contrast, combining multiple models is an active area that excites considerable interest at present because the various methods such as stacking, bagging, and boosting bring into question some of the basic heuristic assumptions about inductive learning. Perhaps Mitchell decided to exclude this subarea because it is in turmoil. Clustering is also not covered, perhaps because it is the focus of a relatively small group within the machine-learning community.

In terms of preparing students for research, students also need to spend time on reading and evaluating current research papers. Students need to learn what constitutes publishable results; what the standards are for various publications; and what constitutes a good research paper, in both form and content. This type of material is outside the scope of any text.

In summary, the text is a pleasure to read. Mitchell's writing style is clear, authoritative, and informative. The detail is not overwhelming, yet a good understanding of the basic concepts and the future of machine learning is provided. The text provides a strong framework for the field of machine learning on which students can build their knowledge.

Dennis Kibler is a professor of information and computer science at the University of California at Irvine. His research has been in machine learning, particularly in problem solving and classification. His e-mail address is kibler@sparky.ics.uci.edu.