FOUNDATIONS OF KNOWLEDGE ACQUISITION: Machine Learning

THE KLUWER INTERNATIONAL SERIES IN ENGINEERING AND COMPUTER SCIENCE

OFFICE OF NAVAL RESEARCH
Advanced Book Series

Consulting Editor

André M. van Tilborg

Other titles in the series:

Foundations of Knowledge Acquisition: Cognitive Models of Complex Learning, edited by Susan Chipman and Alan L. Meyrowitz
ISBN: 0-7923-9277-9

Foundations of Real-Time Computing: Formal Specifications and Methods, edited by André M. van Tilborg and Gary M. Koob
ISBN: 0-7923-9167-5

Foundations of Real-Time Computing: Scheduling and Resource Management, edited by André M. van Tilborg and Gary M. Koob
ISBN: 0-7923-9166-7

FOUNDATIONS OF KNOWLEDGE ACQUISITION: Machine Learning

edited
by

Alan L. Meyrowitz
Naval Research Laboratory

Susan Chipman
Office of Naval Research

KLUWER ACADEMIC PUBLISHERS
Boston/Dordrecht/London

Distributors for North America:
Kluwer Academic Publishers
101 Philip Drive
Assinippi Park
Norwell, Massachusetts 02061 USA

Distributors for all other countries:
Kluwer Academic Publishers Group
Distribution Centre
Post Office Box 322
3300 AH Dordrecht, THE NETHERLANDS

Library of Congress Cataloging-in-Publication Data
(Revised for vol. 2)

Foundations of knowledge acquisition.

(The Kluwer international series in engineering and computer science ; SECS 194)
Editors' names reversed in v. 2.
Includes bibliographical references and index.
Contents: v. [1] Cognitive models of complex learning — v. [2] Machine learning.
1. Knowledge acquisition (Expert systems) I. Chipman, Susan. II. Meyrowitz, Alan Lester. III. Series.
QA76.E95F68 1993 006.3'1 92-36720
ISBN 0-7923-9277-9 (v. 1)
ISBN 0-7923-9278-7 (v. 2)

Printed on acid-free paper.

Printed in the United States of America

TABLE OF CONTENTS

Foreword

One of the most intriguing questions about the new computer technology that has appeared over the past few decades is whether we humans will ever be able to make computers learn. As is painfully obvious to even the most casual computer user, most current computers do not. Yet if we could devise learning techniques that enable computers to routinely improve their performance through experience, the impact would be enormous. The result would be an explosion of new computer applications that would suddenly become economically feasible (e.g., personalized computer assistants that automatically tune themselves to the needs of individual users), and a dramatic improvement in the quality of current computer applications (e.g., imagine an airline scheduling program that improves its scheduling method based on analyzing past delays). And while the potential economic impact of successful learning methods is sufficient reason to invest in research into machine learning, there is a second significant reason: studying machine learning helps us understand our own human learning abilities and disabilities, leading to the possibility of improved methods in education.

While many open questions remain about the methods by which machines and humans might learn, significant progress has been made. For example, learning systems have been demonstrated for tasks such as learning how to drive a vehicle along a roadway (one has successfully driven at 55 mph for 20 miles on a public highway), for learning to evaluate financial loan applications (such systems are now in commercial use), and for learning to recognize human speech (today's top speech recognition systems all employ learning methods). At the same time, a theoretical understanding of learning has begun to appear. For example, we now can place theoretical bounds on the amount of training data a learner must observe in order to reduce its risk of choosing an incorrect hypothesis below some desired threshold. And an improved understanding of human learning is beginning to emerge alongside our improved understanding of machine learning. For example, we now have models of how human novices learn to become experts at various tasks -- models that have been implemented as precise computer programs, and that generate traces very much like those observed in human protocols.

The book you are holding describes a variety of these new results. This work has been pursued under research funding from the Office of Naval Research (ONR) during the time that the editors of this book managed an Accelerated Research Initiative in this area. While several government and private organizations have been important in supporting machine learning research, this ONR effort stands out in particular for its farsighted vision in selecting research topics. During a period when much funding for basic research was being rechanneled to shorter-term development and demonstration projects, ONR had the vision to continue its tradition of supporting research of fundamental long-range significance. The results represent real progress on central problems of machine learning. I encourage you to explore them for yourself in the following chapters.

Tom Mitchell
Carnegie Mellon University

Preface

The two volumes of *Foundations of Knowledge Acquisition* document the recent progress of basic research in knowledge acquisition sponsored by the Office of Naval Research. This volume you are holding is subtitled: *Machine Learning*, and there is a companion volume subtitled: *Cognitive Models of Complex Learning*. Funding was provided by a five-year Accelerated Research Initiative (ARI) from 1988 through 1992, and made possible significant advances in the scientific understanding of how machines and humans can acquire new knowledge so as to exhibit improved problem-solving behavior.

Previous research in artificial intelligence had been directed at understanding the automation of reasoning required for problem solving in complex domains; consequent advances in expert system technology attest to the progress made in the area of deductive inference. However, that research also suggested that automated reasoning can serve to do more than solve a given problem. It can be utilized to infer new facts likely to be useful in tackling future problems, and it can aid in creating new problem-solving strategies. Research sponsored by the Knowledge Acquisition ARI was thus motivated by a desire to understand those reasoning processes which account for the ability of intelligent systems to learn and so improve their performance over time. Such processes can take a variety of forms, including generalization of current knowledge by induction, reasoning by analogy, and discovery (heuristically guided deduction which proceeds from first principles, or axioms). Associated with each are issues regarding the appropriate representation of knowledge to facilitate learning, and the nature of strategies appropriate for learning different kinds of knowledge in diverse domains. There are also issues of computational complexity related to theoretical bounds on what these forms of reasoning can accomplish.

Significant progress in machine learning is reported along a variety of fronts. Chapters in *Machine Learning* include work in analogical reasoning; induction and discovery; learning and planning; learning by competition, using genetic algorithms; and theoretical limitations.

Knowledge acquisition, as pursued under the ARI, was a coordinated research thrust into both machine learning and human learning. Chapters in the companion volume, *Cognitive Models of Complex Learning,* also published by Kluwer Academic Publishers, include summaries of work by cognitive scientists who do computational modeling of human learning. In fact, an accomplishment of research previously sponsored by ONR's Cognitive Science Program was insight into the knowledge and skills that distinguish human novices from human experts in various domains; the Cognitive interest in the ARI was then to characterize how the transition from novice to expert actually takes place. Chapters particularly relevant to that concern are those written by Anderson, Kieras, Marshall, Ohlsson, and VanLehn.

The editors believe these to be valuable volumes from a number of perspectives. They bring together descriptions of recent and on-going research by scientists at the forefront of progress in one of the most challenging arenas of artificial intelligence and cognitive science. Moreover, those scientists were asked to comment on exciting future directions for research in their specialties, and were encouraged to reflect on the progress of science which might go beyond the confines of their particular projects.

Dr. Alan L. Meyrowitz
Navy Center for Applied Research
in Artificial Intelligence

Dr. Susan Chipman
ONR Cognitive Science Program

FOUNDATIONS OF KNOWLEDGE ACQUISITION: Machine Learning