

SYMBOLIC  
COMPUTATION

# Machine Learning

An Artificial Intelligence Approach

Edited by

Ryszard S. Michalski Jaime G. Carbonell

Tom M. Mitchell



Springer-Verlag

Berlin Heidelberg New York Tokyo



C1022

# Machine Learning

## An Artificial Intelligence Approach

Edited by

R. S. Michalski J. G. Carbonell T. M. Mitchell

With Contributions by

J. Anderson R. Banerji G. Bradshaw J. Carbonell  
T. Dietterich N. Haas F. Hayes-Roth G. Hendrix  
P. Langley D. Lenat R. Michalski T. Mitchell J. Mostow  
B. Nudel M. Rychener R. Quinlan H. Simon  
D. Sleeman R. Stepp P. Utgoff



Springer-Verlag  
Berlin Heidelberg New York Tokyo 1984

BIBLIOTHEQUE DU CERIST

## PREFACE

The ability to learn is one of the most fundamental attributes of intelligent behavior. Consequently, progress in the theory and computer modeling of learning processes is of great significance to fields concerned with understanding intelligence. Such fields include cognitive science, artificial intelligence, information science, pattern recognition, psychology, education, epistemology, philosophy, and related disciplines.

The recent observance of the silver anniversary of artificial intelligence has been heralded by a surge of interest in machine learning—both in building models of human learning and in understanding how machines might be endowed with the ability to learn. This renewed interest has spawned many new research projects and resulted in an increase in related scientific activities. In the summer of 1980, the First Machine Learning Workshop was held at Carnegie-Mellon University in Pittsburgh. In the same year, three consecutive issues of the *International Journal of Policy Analysis and Information Systems* were specially devoted to machine learning (No. 2, 3 and 4, 1980). In the spring of 1981, a special issue of the *SIGART Newsletter* No. 76 reviewed current research projects in the field.

This book contains tutorial overviews and research papers representative of contemporary trends in the area of machine learning as viewed from an artificial intelligence perspective. As the first available text on this subject, it is intended to fulfill several needs. For researchers in artificial intelligence, computer science, and cognitive psychology, it provides an easily accessible collection of state-of-the-art papers presenting current results, which will hopefully spur further research. For students in artificial intelligence and related disciplines, this volume may serve as a supplementary textbook for a course in artificial intelligence, or as a primary text for a specialized course devoted to machine learning. Finally, due to the potential impact of machine learning on a variety of disciplines, this book may be of interest to a diverse range of readers, including computer scientists, robotics experts, knowledge engineers, educators, philosophers, data analysts, psychologists and electronic engineers.

The major contemporary research directions in machine learning covered in this book include: learning from examples, modeling human learning strategies, knowledge acquisition for expert systems, learning heuristics, learning from instruction, learning by analogy, discovery systems, and conceptual data analysis. A glossary of selected terminology and an extensive up-to-date bibliography are provided to facilitate instruction and suggest further reading.

*—Ryszard S. Michalski, Jaime G. Carbonell, Tom M. Mitchell*

## ACKNOWLEDGMENTS

The editors wish to express a deep gratitude to all the authors whose efforts made this book possible, and to the reviewers whose ideas and criticism were indispensable in refereeing and improving the contributions. They are most grateful to the Office of Naval Research for supporting the First Machine Learning Workshop, where the idea of publishing this book originated.

The following people provided invaluable assistance in preparing the book: Margorie Ast, Trina and Nathaniel Borenstein, Kitty Fischer, Edward Frank, Jo Ann Gabinelli, Joan Mitchell, June Wingler and Jan Zubkoff. Their hard work and perseverance are gratefully acknowledged. The editors also acknowledge the support and access to technical facilities provided by the Computer Science Departments of Carnegie-Mellon University, the University of Illinois at Urbana-Champaign, and Rutgers University.



# CONTENTS

<b>Preface</b>		<b>v</b>
<b>PART ONE</b>	<b>GENERAL ISSUES IN MACHINE LEARNING</b>	<b>1</b>
<b>Chapter 1</b>	<b>An Overview of Machine Learning</b>	<b>3</b>
	<i>Jaime G. Carbonell, Ryszard S. Michalski, and Tom M. Mitchell</i>	
1.1	Introduction	3
1.2	The Objectives of Machine Learning	3
1.3	A Taxonomy of Machine Learning Research	7
1.4	An Historical Sketch of Machine Learning	14
1.5	A Brief Reader's Guide	16
<b>Chapter 2</b>	<b>Why Should Machines Learn?</b>	<b>25</b>
	<i>Herbert A. Simon</i>	
2.1	Introduction	25
2.2	Human Learning and Machine Learning	25
2.3	What is Learning?	28
2.4	Some Learning Programs	30
2.5	Growth of Knowledge in Large Systems	32
2.6	A Role for Learning	34
2.7	Concluding Remarks	35
<b>PART TWO</b>	<b>LEARNING FROM EXAMPLES</b>	<b>39</b>
<b>Chapter 3</b>	<b>A Comparative Review of Selected Methods for Learning from Examples</b>	<b>41</b>
	<i>Thomas G. Dietterich and Ryszard S. Michalski</i>	
3.1	Introduction	41
3.2	Comparative Review of Selected Methods	49

3.3	Conclusion	75
<b>Chapter 4</b>	<b>A Theory and Methodology of Inductive Learning</b>	<b>83</b>
	<i>Ryszard S. Michalski</i>	
4.1	Introduction	83
4.2	Types of Inductive Learning	87
4.3	Description Language	94
4.4	Problem Background Knowledge	96
4.5	Generalization Rules	103
4.6	The Star Methodology	112
4.7	An Example	116
4.8	Conclusion	123
4.A	Annotated Predicate Calculus (APC)	130
<b>PART THREE</b>	<b>LEARNING IN PROBLEM-SOLVING AND PLANNING</b>	<b>135</b>
<b>Chapter 5</b>	<b>Learning by Analogy: Formulating and Generalizing Plans from Past Experience</b>	<b>137</b>
	<i>Jaime G. Carbonell</i>	
5.1	Introduction	137
5.2	Problem-Solving by Analogy	139
5.3	Evaluating the Analogical Reasoning Process	149
5.4	Learning Generalized Plans	151
5.5	Concluding Remark	159
<b>Chapter 6</b>	<b>Learning by Experimentation: Acquiring and Refining Problem-Solving Heuristics</b>	<b>163</b>
	<i>Tom M. Mitchell, Paul E. Utgoff, and Ranan Banerji</i>	
6.1	Introduction	163
6.2	The Problem	164
6.3	Design of LEX	167
6.4	New Directions: Adding Knowledge to Augment Learning	180
6.5	Summary	189
<b>Chapter 7</b>	<b>Acquisition of Proof Skills in Geometry</b>	<b>191</b>
	<i>John R. Anderson</i>	
7.1	Introduction	191
7.2	A Model of the Skill Underlying Proof Generation	193
7.3	Learning	201
7.4	Knowledge Compilation	202

7.5	Summary of Geometry Learning	217
<b>Chapter 8</b>	<b>Using Proofs and Refutations to Learn from Experience</b>	<b>221</b>
	<i>Frederick Hayes-Roth</i>	
8.1	Introduction	221
8.2	The Learning Cycle	222
8.3	Five Heuristics for Rectifying Refuted Theories	225
8.4	Computational Problems and Implementation Techniques	234
8.5	Conclusions	238
<b>PART FOUR</b>	<b>LEARNING FROM OBSERVATION AND DISCOVERY</b>	<b>241</b>
<b>Chapter 9</b>	<b>The Role of Heuristics in Learning by Discovery: Three Case Studies</b>	<b>243</b>
	<i>Douglas B. Lenat</i>	
9.1	Motivation	243
9.2	Overview	245
9.3	Case Study 1: The AM Program; Heuristics Used to Develop New Knowledge	249
9.4	A Theory of Heuristics	263
9.5	Case Study 2: The Eurisko Program; Heuristics Used to Develop New Heuristics	276
9.6	Heuristics Used to Develop New Representations	282
9.7	Case Study 3: Biological Evolution; Heuristics Used to Generate Plausible Mutations	286
9.8	Conclusions	302
<b>Chapter 10</b>	<b>Rediscovering Chemistry With the BACON System</b>	<b>307</b>
	<i>Pat Langley, Gary L. Bradshaw, and Herbert A. Simon</i>	
10.1	Introduction	307
10.2	An Overview of BACON.4	309
10.3	The Discoveries of BACON.4	312
10.4	Rediscovering Nineteenth Century Chemistry	319
10.5	Conclusions	326



<b>Chapter 11</b>	<b>Learning From Observation: Conceptual Clustering</b>	<b>331</b>
	<i>Ryszard S. Michalski and Robert E. Stepp</i>	
11.1	Introduction	332
11.2	Conceptual Cohesiveness	333
11.3	Terminology and Basic Operations of the Algorithm	336
11.4	A Criterion of Clustering Quality	344
11.5	Method and Implementation	345
11.6	An Example of a Practical Problem: Constructing a Classification Hierarchy of Spanish Folk Songs	358
11.7	Summary and Some Suggested Extensions of the Method	360
<b>PART FIVE</b>	<b>LEARNING FROM INSTRUCTION</b>	<b>365</b>
<b>Chapter 12</b>	<b>Machine Transformation of Advice into a Heuristic Search Procedure</b>	<b>367</b>
	<i>David Jack Mostow</i>	
12.1	Introduction	367
12.2	Kinds of Knowledge Used	370
12.3	A Slightly Non-Standard Definition of Heuristic Search	374
12.4	Instantiating the HSM Schema for a Given Problem	378
12.5	Refining HSM by Moving Constraints Between Control Components	384
12.6	Evaluation of Generality	398
12.7	Conclusion	399
12.A	Index of Rules	403
<b>Chapter 13</b>	<b>Learning by Being Told: Acquiring Knowledge for Information Management</b>	<b>405</b>
	<i>Norm Haas and Gary G. Hendrix</i>	
13.1	Overview	405
13.2	Technical Approach: Experiments with the KLAUS Concept	408
13.3	More Technical Details	413
13.4	Conclusions and Directions for Future Work	418
13.A	Training NANOKLAUS About Aircraft Carriers	422

<b>Chapter 14</b>	<b>The Instructible Production System: A Retrospective Analysis</b>	<b>429</b>
	<i>Michael D. Rychener</i>	
14.1	The Instructible Production System Project	430
14.2	Essential Functional Components of Instructible Systems	436
14.3	Survey of Approaches	443
14.4	Discussion	453
<b>PART SIX</b>	<b>APPLIED LEARNING SYSTEMS</b>	<b>461</b>
<b>Chapter 15</b>	<b>Learning Efficient Classification Procedures and their Application to Chess End Games</b>	<b>463</b>
	<i>J. Ross Quinlan</i>	
15.1	Introduction	463
15.2	The Inductive Inference Machinery	465
15.3	The Lost N-ply Experiments	470
15.4	Approximate Classification Rules	474
15.5	Some Thoughts on Discovering Attributes	477
15.6	Conclusion	481
<b>Chapter 16</b>	<b>Inferring Student Models for Intelligent Computer-Aided Instruction</b>	<b>483</b>
	<i>Derek H. Sleeman</i>	
16.1	Introduction	483
16.2	Generating a Complete and Non-redundant Set of Models	488
16.3	Processing Domain Knowledge	503
16.4	Summary	507
16.A	An Example of the SELECTIVE Algorithm: LMS-I's Model Generation Algorithm	510
<b>Comprehensive Bibliography of Machine Learning</b>		<b>511</b>
	<i>Paul E. Utgoff and Bernard Nudel</i>	
<b>Glossary of Selected Terms In Machine Learning</b>		<b>551</b>
<b>About the Authors</b>		<b>557</b>
<b>Author Index</b>		<b>563</b>
<b>Subject Index</b>		<b>567</b>