DATA MINING WITH DECISION TREES

Theory and Applications

2nd Edition

SERIES IN MACHINE PERCEPTION AND ARTIFICIAL INTELLIGENCE*

Editors: **H. Bunke** (Univ. Bern, Switzerland)
P. S. P. Wang (Northeastern Univ., USA)

- Vol. 65: Fighting Terror in Cyberspace
(Eds. *M. Last and A. Kandel*)
- Vol. 66: Formal Models, Languages and Applications
(Eds. *K. G. Subramanian, K. Rangarajan and M. Mukund*)
- Vol. 67: Image Pattern Recognition: Synthesis and Analysis in Biometrics
(Eds. *S. N. Yanushkevich, P. S. P. Wang, M. L. Gavrilova and S. N. Srihari*)
- Vol. 68: Bridging the Gap Between Graph Edit Distance and Kernel Machines
(*M. Neuhaus and H. Bunke*)
- Vol. 69: Data Mining with Decision Trees: Theory and Applications
(*L. Rokach and O. Maimon*)
- Vol. 70: Personalization Techniques and Recommender Systems
(Eds. *G. Uchytel and M. Ma*)
- Vol. 71: Recognition of Whiteboard Notes: Online, Offline and Combination
(Eds. *H. Bunke and M. Liwicki*)
- Vol. 72: Kernels for Structured Data
(*T Gärtner*)
- Vol. 73: Progress in Computer Vision and Image Analysis
(Eds. *H. Bunke, J. J. Villanueva, G. Sánchez and X. Otazu*)
- Vol. 74: Wavelet Theory Approach to Pattern Recognition (2nd Edition)
(*Y. Y. Tang*)
- Vol. 75: Pattern Classification Using Ensemble Methods
(*L. Rokach*)
- Vol. 76: Automated Database Applications Testing: Specification Representation for Automated Reasoning
(*R. F. Mikhail, D. Berndt and A. Kandel*)
- Vol. 77: Graph Classification and Clustering Based on Vector Space Embedding
(*K. Riesen and H. Bunke*)
- Vol. 78: Integration of Swarm Intelligence and Artificial Neural Network
(Eds. *S. Dehuri, S. Ghosh and S.-B. Cho*)
- Vol. 79: Document Analysis and Recognition with Wavelet and Fractal Theories
(*Y. Y. Tang*)
- Vol. 80: Multimodal Interactive Handwritten Text Transcription
(*V. Romero, A. H. Toselli and E. Vidal*)
- Vol. 81: Data Mining with Decision Trees: Theory and Applications
Second Edition
(*L. Rokach and O. Maimon*)

*The complete list of the published volumes in the series can be found at
<http://www.worldscientific.com/series/smpai>

DATA MINING WITH DECISION TREES

Theory and Applications

2nd Edition

Lior Rokach

Ben-Gurion University of the Negev, Israel

Oded Maimon

Tel-Aviv University, Israel

Published by

World Scientific Publishing Co. Pte. Ltd.

5 Toh Tuck Link, Singapore 596224

USA office: 27 Warren Street, Suite 401-402, Hackensack, NJ 07601

UK office: 57 Shelton Street, Covent Garden, London WC2H 9HE

Library of Congress Cataloging-in-Publication Data

Rokach, Lior.

Data mining with decision trees : theory and applications / by Lior Rokach (Ben-Gurion University of the Negev, Israel), Oded Maimon (Tel-Aviv University, Israel). -- 2nd edition.
pages cm

Includes bibliographical references and index.

ISBN 978-9814590075 (hardback : alk. paper) -- ISBN 978-9814590082 (ebook)

1. Data mining. 2. Decision trees. 3. Machine learning. 4. Decision support systems.

I. Maimon, Oded. II. Title.

QA76.9.D343R654 2014

006.3'12--dc23

2014029799

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library.

Copyright © 2015 by World Scientific Publishing Co. Pte. Ltd.

All rights reserved. This book, or parts thereof, may not be reproduced in any form or by any means, electronic or mechanical, including photocopying, recording or any information storage and retrieval system now known or to be invented, without written permission from the publisher.

For photocopying of material in this volume, please pay a copying fee through the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, USA. In this case permission to photocopy is not required from the publisher.

In-house Editor: Amanda Yun

Typeset by Stallion Press

Email: enquiries@stallionpress.com

Printed in Singapore

Dedicated to our families
in appreciation for their patience and support
during the preparation of this book.

L.R.
O.M.

About the Authors

Lior Rokach is an Associate Professor of Information Systems and Software Engineering at Ben-Gurion University of the Negev. Dr. Rokach is a recognized expert in intelligent information systems and has held several leading positions in this field. His main areas of interest are Machine Learning, Information Security, Recommender Systems and Information Retrieval. Dr. Rokach is the author of over 100 peer reviewed papers in leading journals conference proceedings, patents, and book chapters. In addition, he has also authored six books in the field of data mining.

Professor **Oded Maimon** from Tel Aviv University, previously at MIT, is also the Oracle chair professor. His research interests are in data mining and knowledge discovery and robotics. He has published over 300 papers and ten books. Currently he is exploring new concepts of core data mining methods, as well as investigating artificial and biological data.

Preface for the Second Edition

The first edition of the book, which was published six years ago, was extremely well received by the data mining research and development communities. The positive reception, along with the fast pace of research in the data mining, motivated us to update our book. We received many requests to include the new advances in the field as well as the new applications and software tools that have become available in the second edition of the book. This second edition aims to refresh the previously presented material in the fundamental areas, and to present new findings in the field; nearly quarter of this edition is comprised of new materials.

We have added four new chapters and updated some of the existing ones. Because many readers are already familiar with the layout of the first edition, we have tried to change it as little as possible. Below is the summary of the main alterations:

- The first edition has mainly focused on using decision trees for classification tasks (i.e. classification trees). In this edition we describe how decision trees can be used for other data mining tasks, such as regression, clustering and survival analysis.
- The new addition includes a walk-through-guide for using decision trees software. Specifically, we focus on open-source solutions that are freely available.
- We added a chapter on cost-sensitive active and proactive learning of decision trees since the cost aspect is very important in many domain applications such as medicine and marketing.
- Chapter 16 is dedicated entirely to the field of recommender systems which is a popular research area. Recommender Systems help customers

to choose an item from a potentially overwhelming number of alternative items.

We apologize for the errors that have been found in the first edition and we are grateful to the many readers who have found those. We have done our best to avoid errors in this new edition. Many graduate students have read parts of the manuscript and offered helpful suggestions and we thank them for that.

Many thanks are owed to Elizaveta Futerman. She has been the most helpful assistant in proofreading the new chapters and improving the manuscript. The authors would like to thank Amanda Yun and staff members of World Scientific Publishing for their kind cooperation in writing this book. Moreover, we are thankful to Prof. H. Bunke and Prof. P.S.P. Wang for including our book in their fascinating series on machine perception and artificial intelligence.

Finally, we would like to thank our families for their love and support.

Beer-Sheva, Israel
Tel-Aviv, Israel

Lior Rokach
Oded Maimon

April 2014

Preface for the First Edition

Data mining is the science, art and technology of exploring large and complex bodies of data in order to discover useful patterns. Theoreticians and practitioners are continually seeking improved techniques to make the process more efficient, cost-effective and accurate. One of the most promising and popular approaches is the use of decision trees. Decision trees are simple yet successful techniques for predicting and explaining the relationship between some measurements about an item and its target value. In addition to their use in data mining, decision trees, which originally derived from logic, management and statistics, are today highly effective tools in other areas such as text mining, information extraction, machine learning, and pattern recognition.

Decision trees offer many benefits:

- Versatility for a wide variety of data mining tasks, such as classification, regression, clustering and feature selection
- Self-explanatory and easy to follow (when compacted)
- Flexibility in handling a variety of input data: nominal, numeric and textual
- Adaptability in processing datasets that may have errors or missing values
- High predictive performance for a relatively small computational effort
- Available in many data mining packages over a variety of platforms
- Useful for large datasets (in an ensemble framework)

This is the first comprehensive book about decision trees. Devoted entirely to the field, it covers almost all aspects of this very important technique.

The book has three main parts:

- Part I presents the data mining and decision tree foundations (including basic rationale, theoretical formulation, and detailed evaluation).
- Part II introduces the basic and advanced algorithms for automatically growing decision trees (including splitting and pruning, decision forests, and incremental learning).
- Part III presents important extensions for improving decision tree performance and for accommodating it to certain circumstances. This part also discusses advanced topics such as feature selection, fuzzy decision trees and hybrid framework.

We have tried to make as complete a presentation of decision trees in data mining as possible. However, new applications are always being introduced. For example, we are now researching the important issue of data mining privacy, where we use a hybrid method of genetic process with decision trees to generate the optimal privacy-protecting method. Using the fundamental techniques presented in this book, we are also extensively involved in researching language-independent text mining (including ontology generation and automatic taxonomy).

Although we discuss in this book the broad range of decision trees and their importance, we are certainly aware of related methods, some with overlapping capabilities. For this reason, we recently published a complementary book “Soft Computing for Knowledge Discovery and Data Mining”, which addresses other approaches and methods in data mining, such as artificial neural networks, fuzzy logic, evolutionary algorithms, agent technology, swarm intelligence and diffusion methods.

An important principle that guided us while writing this book was the extensive use of illustrative examples. Accordingly, in addition to decision tree theory and algorithms, we provide the reader with many applications from the real-world as well as examples that we have formulated for explaining the theory and algorithms. The applications cover a variety of fields, such as marketing, manufacturing, and bio-medicine. The data referred to in this book, as well as most of the Java implementations of the pseudo-algorithms and programs that we present and discuss, may be obtained via the Web.

We believe that this book will serve as a vital source of decision tree techniques for researchers in information systems, engineering, computer science, statistics and management. In addition, this book is highly useful to researchers in the social sciences, psychology, medicine, genetics, business

intelligence, and other fields characterized by complex data-processing problems of underlying models.

Since the material in this book formed the basis of undergraduate and graduates courses at Ben-Gurion University of the Negev and Tel-Aviv University and it can also serve as a reference source for graduate/advanced undergraduate level courses in knowledge discovery, data mining and machine learning. Practitioners among the readers may be particularly interested in the descriptions of real-world data mining projects performed with decision trees methods.

We would like to acknowledge the contribution to our research and to the book to many students, but in particular to Dr. Barak Chizi, Dr. Shahar Cohen, Roni Romano and Reuven Arbel. Many thanks are owed to Arthur Kemelman. He has been a most helpful assistant in proofreading and improving the manuscript.

The authors would like to thank Mr. Ian Seldrup, Senior Editor, and staff members of World Scientific Publishing for their kind cooperation in connection with writing this book. Thanks also to Prof. H. Bunke and Prof P.S.P. Wang for including our book in their fascinating series in machine perception and artificial intelligence.

Last, but not least, we owe our special gratitude to our partners, families, and friends for their patience, time, support, and encouragement.

Beer-Sheva, Israel
Tel-Aviv, Israel

Lior Rokach
Oded Maimon

October 2007

This page intentionally left blank

Contents

<i>About the Authors</i>	vi
<i>Preface for the Second Edition</i>	vii
<i>Preface for the First Edition</i>	ix
1. Introduction to Decision Trees	1
1.1 Data Science	1
1.2 Data Mining	2
1.3 The Four-Layer Model	3
1.4 Knowledge Discovery in Databases (KDD)	4
1.5 Taxonomy of Data Mining Methods	8
1.6 Supervised Methods	9
1.6.1 Overview	9
1.7 Classification Trees	10
1.8 Characteristics of Classification Trees	12
1.8.1 Tree Size	14
1.8.2 The Hierarchical Nature of Decision Trees	15
1.9 Relation to Rule Induction	15
2. Training Decision Trees	17
2.1 What is Learning?	17
2.2 Preparing the Training Set	17
2.3 Training the Decision Tree	19

- 3. A Generic Algorithm for Top-Down Induction of Decision Trees 23
 - 3.1 Training Set 23
 - 3.2 Definition of the Classification Problem 25
 - 3.3 Induction Algorithms 26
 - 3.4 Probability Estimation in Decision Trees 26
 - 3.4.1 Laplace Correction 27
 - 3.4.2 No Match 28
 - 3.5 Algorithmic Framework for Decision Trees 28
 - 3.6 Stopping Criteria 30
- 4. Evaluation of Classification Trees 31
 - 4.1 Overview 31
 - 4.2 Generalization Error 31
 - 4.2.1 Theoretical Estimation of Generalization Error 32
 - 4.2.2 Empirical Estimation of Generalization Error 32
 - 4.2.3 Alternatives to the Accuracy Measure 34
 - 4.2.4 The F-Measure 35
 - 4.2.5 Confusion Matrix 36
 - 4.2.6 Classifier Evaluation under Limited Resources 37
 - 4.2.6.1 ROC Curves 39
 - 4.2.6.2 Hit-Rate Curve 40
 - 4.2.6.3 Qrecall (Quota Recall) 40
 - 4.2.6.4 Lift Curve 41
 - 4.2.6.5 Pearson Correlation Coefficient 41
 - 4.2.6.6 Area Under Curve (AUC) 43
 - 4.2.6.7 Average Hit-Rate 44
 - 4.2.6.8 Average Qrecall 44
 - 4.2.6.9 Potential Extract Measure (PEM) 45
 - 4.2.7 Which Decision Tree Classifier is Better? 48
 - 4.2.7.1 McNemar’s Test 48
 - 4.2.7.2 A Test for the Difference of Two Proportions 50
 - 4.2.7.3 The Resampled Paired t Test 51
 - 4.2.7.4 The k -fold Cross-validated Paired t Test 51
 - 4.3 Computational Complexity 52

4.4	Comprehensibility	52
4.5	Scalability to Large Datasets	53
4.6	Robustness	55
4.7	Stability	55
4.8	Interestingness Measures	56
4.9	Overfitting and Underfitting	57
4.10	“No Free Lunch” Theorem	58
5.	Splitting Criteria	61
5.1	Univariate Splitting Criteria	61
5.1.1	Overview	61
5.1.2	Impurity-based Criteria	61
5.1.3	Information Gain	62
5.1.4	Gini Index	62
5.1.5	Likelihood Ratio Chi-squared Statistics	63
5.1.6	DKM Criterion	63
5.1.7	Normalized Impurity-based Criteria	63
5.1.8	Gain Ratio	64
5.1.9	Distance Measure	64
5.1.10	Binary Criteria	64
5.1.11	Twoing Criterion	65
5.1.12	Orthogonal Criterion	65
5.1.13	Kolmogorov–Smirnov Criterion	66
5.1.14	AUC Splitting Criteria	66
5.1.15	Other Univariate Splitting Criteria	66
5.1.16	Comparison of Univariate Splitting Criteria	66
5.2	Handling Missing Values	67
6.	Pruning Trees	69
6.1	Stopping Criteria	69
6.2	Heuristic Pruning	69
6.2.1	Overview	69
6.2.2	Cost Complexity Pruning	70
6.2.3	Reduced Error Pruning	70
6.2.4	Minimum Error Pruning (MEP)	71
6.2.5	Pessimistic Pruning	71
6.2.6	Error-Based Pruning (EBP)	72
6.2.7	Minimum Description Length (MDL) Pruning	73
6.2.8	Other Pruning Methods	73

- 6.2.9 Comparison of Pruning Methods 73
- 6.3 Optimal Pruning 74
- 7. Popular Decision Trees Induction Algorithms 77
 - 7.1 Overview 77
 - 7.2 ID3 77
 - 7.3 C4.5 78
 - 7.4 CART 79
 - 7.5 CHAID 79
 - 7.6 QUEST 80
 - 7.7 Reference to Other Algorithms 80
 - 7.8 Advantages and Disadvantages of Decision Trees 81
- 8. Beyond Classification Tasks 85
 - 8.1 Introduction 85
 - 8.2 Regression Trees 85
 - 8.3 Survival Trees 86
 - 8.4 Clustering Tree 89
 - 8.4.1 Distance Measures 89
 - 8.4.2 Minkowski: Distance Measures for Numeric Attributes 90
 - 8.4.2.1 Distance Measures for Binary Attributes 90
 - 8.4.2.2 Distance Measures for Nominal Attributes 91
 - 8.4.2.3 Distance Metrics for Ordinal Attributes 91
 - 8.4.2.4 Distance Metrics for Mixed-Type Attributes 92
 - 8.4.3 Similarity Functions 92
 - 8.4.3.1 Cosine Measure 93
 - 8.4.3.2 Pearson Correlation Measure 93
 - 8.4.3.3 Extended Jaccard Measure 93
 - 8.4.3.4 Dice Coefficient Measure 93
 - 8.4.4 The OCCT Algorithm 93
 - 8.5 Hidden Markov Model Trees 94
- 9. Decision Forests 99
 - 9.1 Introduction 99
 - 9.2 Back to the Roots 99

9.3	Combination Methods	108
9.3.1	Weighting Methods	108
9.3.1.1	Majority Voting	108
9.3.1.2	Performance Weighting	109
9.3.1.3	Distribution Summation	109
9.3.1.4	Bayesian Combination	109
9.3.1.5	Dempster–Shafer	110
9.3.1.6	Vogging	110
9.3.1.7	Naïve Bayes	110
9.3.1.8	Entropy Weighting	110
9.3.1.9	Density-based Weighting	111
9.3.1.10	DEA Weighting Method	111
9.3.1.11	Logarithmic Opinion Pool	111
9.3.1.12	Gating Network	112
9.3.1.13	Order Statistics	113
9.3.2	Meta-combination Methods	113
9.3.2.1	Stacking	113
9.3.2.2	Arbiter Trees	114
9.3.2.3	Combiner Trees	116
9.3.2.4	Grading	117
9.4	Classifier Dependency	118
9.4.1	Dependent Methods	118
9.4.1.1	Model-guided Instance Selection	118
9.4.1.2	Incremental Batch Learning	122
9.4.2	Independent Methods	122
9.4.2.1	Bagging	122
9.4.2.2	Wagging	124
9.4.2.3	Random Forest	125
9.4.2.4	Rotation Forest	126
9.4.2.5	Cross-validated Committees	129
9.5	Ensemble Diversity	130
9.5.1	Manipulating the Inducer	131
9.5.1.1	Manipulation of the Inducer’s Parameters	131
9.5.1.2	Starting Point in Hypothesis Space	132
9.5.1.3	Hypothesis Space Traversal	132
9.5.1.3.1	Random-based Strategy	132
9.5.1.3.2	Collective-Performance-based Strategy	132

9.5.2	Manipulating the Training Samples	133
9.5.2.1	Resampling	133
9.5.2.2	Creation	133
9.5.2.3	Partitioning	134
9.5.3	Manipulating the Target Attribute Representation	134
9.5.4	Partitioning the Search Space	136
9.5.4.1	Divide and Conquer	136
9.5.4.2	Feature Subset-based Ensemble Methods	137
9.5.4.2.1	Random-based Strategy	138
9.5.4.2.2	Reduct-based Strategy	138
9.5.4.2.3	Collective-Performance-based Strategy	139
9.5.4.2.4	Feature Set Partitioning	139
9.5.5	Multi-Inducers	142
9.5.6	Measuring the Diversity	143
9.6	Ensemble Size	144
9.6.1	Selecting the Ensemble Size	144
9.6.2	Pre-selection of the Ensemble Size	145
9.6.3	Selection of the Ensemble Size while Training	145
9.6.4	Pruning — Post Selection of the Ensemble Size	146
9.6.4.1	Pre-combining Pruning	146
9.6.4.2	Post-combining Pruning	146
9.7	Cross-Inducer	147
9.8	Multistrategy Ensemble Learning	148
9.9	Which Ensemble Method Should be Used?	148
9.10	Open Source for Decision Trees Forests	149
10.	A Walk-through-guide for Using Decision Trees Software	151
10.1	Introduction	151
10.2	Weka	152
10.2.1	Training a Classification Tree	153
10.2.2	Building a Forest	158
10.3	R	159
10.3.1	Party Package	159
10.3.2	Forest	162

10.3.3	Other Types of Trees	163
10.3.4	The Rpart Package	164
10.3.5	RandomForest	165
11.	Advanced Decision Trees	167
11.1	Oblivious Decision Trees	167
11.2	Online Adaptive Decision Trees	168
11.3	Lazy Tree	168
11.4	Option Tree	169
11.5	Lookahead	172
11.6	Oblique Decision Trees	172
11.7	Incremental Learning of Decision Trees	175
11.7.1	The Motives for Incremental Learning	175
11.7.2	The Inefficiency Challenge	176
11.7.3	The Concept Drift Challenge	177
11.8	Decision Trees Inducers for Large Datasets	179
11.8.1	Accelerating Tree Induction	180
11.8.2	Parallel Induction of Tree	182
12.	Cost-sensitive Active and Proactive Learning of Decision Trees	183
12.1	Overview	183
12.2	Type of Costs	184
12.3	Learning with Costs	185
12.4	Induction of Cost Sensitive Decision Trees	188
12.5	Active Learning	189
12.6	Proactive Data Mining	196
12.6.1	Changing the Input Data	197
12.6.2	Attribute Changing Cost and Benefit Functions	198
12.6.3	Maximizing Utility	199
12.6.4	An Algorithmic Framework for Proactive Data Mining	200
13.	Feature Selection	203
13.1	Overview	203
13.2	The “Curse of Dimensionality”	203
13.3	Techniques for Feature Selection	206
13.3.1	Feature Filters	207

13.3.1.1	FOCUS	207
13.3.1.2	LVF	207
13.3.1.3	Using a Learning Algorithm as a Filter	207
13.3.1.4	An Information Theoretic Feature Filter	208
13.3.1.5	RELIEF Algorithm	208
13.3.1.6	Simba and G-flip	208
13.3.1.7	Contextual Merit (CM) Algorithm . . .	209
13.3.2	Using Traditional Statistics for Filtering	209
13.3.2.1	Mallows Cp	209
13.3.2.2	AIC, BIC and F-ratio	209
13.3.2.3	Principal Component Analysis (PCA)	210
13.3.2.4	Factor Analysis (FA)	210
13.3.2.5	Projection Pursuit (PP)	210
13.3.3	Wrappers	211
13.3.3.1	Wrappers for Decision Tree Learners	211
13.4	Feature Selection as a means of Creating Ensembles . . .	211
13.5	Ensemble Methodology for Improving Feature Selection	213
13.5.1	Independent Algorithmic Framework	215
13.5.2	Combining Procedure	216
13.5.2.1	Simple Weighted Voting	216
13.5.2.2	Using Artificial Contrasts	218
13.5.3	Feature Ensemble Generator	220
13.5.3.1	Multiple Feature Selectors	220
13.5.3.2	Bagging	221
13.6	Using Decision Trees for Feature Selection	221
13.7	Limitation of Feature Selection Methods	222
14.	Fuzzy Decision Trees	225
14.1	Overview	225
14.2	Membership Function	226
14.3	Fuzzy Classification Problems	227
14.4	Fuzzy Set Operations	228
14.5	Fuzzy Classification Rules	229
14.6	Creating Fuzzy Decision Tree	230

14.6.1	Fuzzifying Numeric Attributes	230
14.6.2	Inducing of Fuzzy Decision Tree	232
14.7	Simplifying the Decision Tree	234
14.8	Classification of New Instances	234
14.9	Other Fuzzy Decision Tree Inducers	234
15.	Hybridization of Decision Trees with other Techniques	237
15.1	Introduction	237
15.2	A Framework for Instance-Space Decomposition	237
15.2.1	Stopping Rules	240
15.2.2	Splitting Rules	241
15.2.3	Split Validation Examinations	241
15.3	The Contrasted Population Miner (CPOM) Algorithm	242
15.3.1	CPOM Outline	242
15.3.2	The Grouped Gain Ratio Splitting Rule	244
15.4	Induction of Decision Trees by an Evolutionary Algorithm (EA)	246
16.	Decision Trees and Recommender Systems	251
16.1	Introduction	251
16.2	Using Decision Trees for Recommending Items	252
16.2.1	RS-Adapted Decision Tree	253
16.2.2	Least Probable Intersections	257
16.3	Using Decision Trees for Preferences Elicitation	259
16.3.1	Static Methods	261
16.3.2	Dynamic Methods and Decision Trees	262
16.3.3	SVD-based CF Method	263
16.3.4	Pairwise Comparisons	264
16.3.5	Profile Representation	266
16.3.6	Selecting the Next Pairwise Comparison	267
16.3.7	Clustering the Items	269
16.3.8	Training a Lazy Decision Tree	270

<i>Bibliography</i>	273
---------------------	-----

<i>Index</i>	303
--------------	-----