

Undergraduate Topics in Computer Science

Undergraduate Topics in Computer Science (UTiCS) delivers high-quality instructional content for undergraduates studying in all areas of computing and information science. From core foundational and theoretical material to final-year topics and applications, UTiCS books take a fresh, concise, and modern approach and are ideal for self-study or for a one- or two-semester course. The texts are all authored by established experts in their fields, reviewed by an international advisory board, and contain numerous examples and problems. Many include fully worked solutions.

Also in this series

Iain Craig
Object-Oriented Programming Languages: Interpretation
978-1-84628-773-2

Hanne Riis Nielson and Flemming Nielson
Semantics with Applications: An Appetizer
978-1-84628-691-9

Max Bramer

Principles of Data Mining

 Springer

Max Bramer, BSc, PhD, CEng, FBCS, FIEE, FRSA
Digital Professor of Information Technology, University of Portsmouth, UK

Series editor

Ian Mackie
École Polytechnique, France and King's College London, UK

Advisory board

Samson Abramsky, University of Oxford, UK
Chris Hankin, Imperial College London, UK
Dexter Kozen, Cornell University, USA
Andrew Pitts, University of Cambridge, UK
Hanne Riis Nielson, Technical University of Denmark, Denmark
Steven Skiena, Stony Brook University, USA
Iain Stewart, University of Durham, UK
David Zhang, The Hong Kong Polytechnic University, Hong Kong

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library

Library of Congress Control Number: 2007922358

Undergraduate Topics in Computer Science ISSN 1863-7310
ISBN-10: 1-84628-765-0 e-ISBN 10: 1-84628-766-9
ISBN-13: 978-1-84628-765-7 e-ISBN-13: 978-1-84628-766-4

Printed on acid-free paper

© Springer-Verlag London Limited 2007

Apart from any fair dealing for the purposes of research or private study, or criticism or review, as permitted under the Copyright, Designs and Patents Act 1988, this publication may only be reproduced, stored or transmitted, in any form or by any means, with the prior permission in writing of the publishers, or in the case of reprographic reproduction in accordance with the terms of licences issued by the Copyright Licensing Agency. Enquiries concerning reproduction outside those terms should be sent to the publishers.

The use of registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant laws and regulations and therefore free for general use.

The publisher makes no representation, express or implied, with regard to the accuracy of the information contained in this book and cannot accept any legal responsibility or liability for any errors or omissions that may be made.

9 8 7 6 5 4 3 2 1

Springer Science+Business Media
springer.com

Contents

Introduction to Data Mining	1
1. Data for Data Mining	11
1.1 Standard Formulation	11
1.2 Types of Variable	12
1.2.1 Categorical and Continuous Attributes	14
1.3 Data Preparation	14
1.3.1 Data Cleaning	15
1.4 Missing Values	17
1.4.1 Discard Instances	17
1.4.2 Replace by Most Frequent/Average Value	17
1.5 Reducing the Number of Attributes	18
1.6 The UCI Repository of Datasets	19
Chapter Summary	20
Self-assessment Exercises for Chapter 1	20
2. Introduction to Classification: Naïve Bayes and Nearest Neighbour	23
2.1 What is Classification?	23
2.2 Naïve Bayes Classifiers	24
2.3 Nearest Neighbour Classification	31
2.3.1 Distance Measures	34
2.3.2 Normalisation	37
2.3.3 Dealing with Categorical Attributes	38
2.4 Eager and Lazy Learning	38
Chapter Summary	39

Self-assessment Exercises for Chapter 2	39
3. Using Decision Trees for Classification	41
3.1 Decision Rules and Decision Trees	41
3.1.1 Decision Trees: The Golf Example	42
3.1.2 Terminology	43
3.1.3 The <i>degrees</i> Dataset	44
3.2 The TDIDT Algorithm.....	47
3.3 Types of Reasoning	49
Chapter Summary	50
Self-assessment Exercises for Chapter 3	50
4. Decision Tree Induction: Using Entropy for Attribute Selection	51
4.1 Attribute Selection: An Experiment.....	51
4.2 Alternative Decision Trees	52
4.2.1 The Football/Netball Example	53
4.2.2 The <i>anonymous</i> Dataset	55
4.3 Choosing Attributes to Split On: Using Entropy	56
4.3.1 The <i>lens24</i> Dataset	57
4.3.2 Entropy	59
4.3.3 Using Entropy for Attribute Selection.....	60
4.3.4 Maximising Information Gain	62
Chapter Summary	63
Self-assessment Exercises for Chapter 4	63
5. Decision Tree Induction: Using Frequency Tables for Attribute Selection	65
5.1 Calculating Entropy in Practice	65
5.1.1 Proof of Equivalence	66
5.1.2 A Note on Zeros	68
5.2 Other Attribute Selection Criteria: Gini Index of Diversity.....	68
5.3 Inductive Bias	70
5.4 Using Gain Ratio for Attribute Selection	72
5.4.1 Properties of Split Information	73
5.5 Number of Rules Generated by Different Attribute Selection Criteria	74
5.6 Missing Branches.....	75
Chapter Summary	76
Self-assessment Exercises for Chapter 5	77

6. Estimating the Predictive Accuracy of a Classifier	79
6.1 Introduction	79
6.2 Method 1: Separate Training and Test Sets	80
6.2.1 Standard Error	81
6.2.2 Repeated Train and Test	82
6.3 Method 2: k -fold Cross-validation	82
6.4 Method 3: N -fold Cross-validation	83
6.5 Experimental Results I	84
6.6 Experimental Results II: Datasets with Missing Values	86
6.6.1 Strategy 1: Discard Instances	87
6.6.2 Strategy 2: Replace by Most Frequent/Average Value	87
6.6.3 Missing Classifications	89
6.7 Confusion Matrix	89
6.7.1 True and False Positives	90
Chapter Summary	91
Self-assessment Exercises for Chapter 6	91
7. Continuous Attributes	93
7.1 Introduction	93
7.2 Local versus Global Discretisation	95
7.3 Adding Local Discretisation to TDIDT	96
7.3.1 Calculating the Information Gain of a Set of Pseudo- attributes	97
7.3.2 Computational Efficiency	102
7.4 Using the ChiMerge Algorithm for Global Discretisation	105
7.4.1 Calculating the Expected Values and χ^2	108
7.4.2 Finding the Threshold Value	113
7.4.3 Setting <i>minIntervals</i> and <i>maxIntervals</i>	113
7.4.4 The ChiMerge Algorithm: Summary	115
7.4.5 The ChiMerge Algorithm: Comments	115
7.5 Comparing Global and Local Discretisation for Tree Induction	116
Chapter Summary	118
Self-assessment Exercises for Chapter 7	118
8. Avoiding Overfitting of Decision Trees	119
8.1 Dealing with Clashes in a Training Set	120
8.1.1 Adapting TDIDT to Deal With Clashes	120
8.2 More About Overfitting Rules to Data	125
8.3 Pre-pruning Decision Trees	126
8.4 Post-pruning Decision Trees	128
Chapter Summary	133
Self-assessment Exercise for Chapter 8	134

9. More About Entropy	135
9.1 Introduction	135
9.2 Coding Information Using Bits	138
9.3 Discriminating Amongst M Values (M Not a Power of 2)	140
9.4 Encoding Values That Are Not Equally Likely	141
9.5 Entropy of a Training Set	144
9.6 Information Gain Must be Positive or Zero	145
9.7 Using Information Gain for Feature Reduction for Classification Tasks	147
9.7.1 Example 1: The <i>genetics</i> Dataset	148
9.7.2 Example 2: The <i>bcst96</i> Dataset	152
Chapter Summary	154
Self-assessment Exercises for Chapter 9	154
10. Inducing Modular Rules for Classification	155
10.1 Rule Post-pruning	155
10.2 Conflict Resolution	157
10.3 Problems with Decision Trees	160
10.4 The Prism Algorithm	162
10.4.1 Changes to the Basic Prism Algorithm	169
10.4.2 Comparing Prism with TDIDT	170
Chapter Summary	171
Self-assessment Exercise for Chapter 10	171
11. Measuring the Performance of a Classifier	173
11.1 True and False Positives and Negatives	174
11.2 Performance Measures	176
11.3 True and False Positive Rates versus Predictive Accuracy	179
11.4 ROC Graphs	180
11.5 ROC Curves	182
11.6 Finding the Best Classifier	183
Chapter Summary	184
Self-assessment Exercise for Chapter 11	185
12. Association Rule Mining I	187
12.1 Introduction	187
12.2 Measures of Rule Interestingness	189
12.2.1 The Piatetsky-Shapiro Criteria and the RI Measure	191
12.2.2 Rule Interestingness Measures Applied to the <i>chess</i> Dataset	193

12.2.3 Using Rule Interestingness Measures for Conflict Resolution	195
12.3 Association Rule Mining Tasks	195
12.4 Finding the Best N Rules	196
12.4.1 The J -Measure: Measuring the Information Content of a Rule	197
12.4.2 Search Strategy	198
Chapter Summary	201
Self-assessment Exercises for Chapter 12	201
13. Association Rule Mining II	203
13.1 Introduction	203
13.2 Transactions and Itemsets	204
13.3 Support for an Itemset	205
13.4 Association Rules	206
13.5 Generating Association Rules	208
13.6 Apriori	209
13.7 Generating Supported Itemsets: An Example	212
13.8 Generating Rules for a Supported Itemset	214
13.9 Rule Interestingness Measures: Lift and Leverage	216
Chapter Summary	218
Self-assessment Exercises for Chapter 13	219
14. Clustering	221
14.1 Introduction	221
14.2 k -Means Clustering	224
14.2.1 Example	225
14.2.2 Finding the Best Set of Clusters	230
14.3 Agglomerative Hierarchical Clustering	231
14.3.1 Recording the Distance Between Clusters	233
14.3.2 Terminating the Clustering Process	236
Chapter Summary	237
Self-assessment Exercises for Chapter 14	238
15. Text Mining	239
15.1 Multiple Classifications	239
15.2 Representing Text Documents for Data Mining	240
15.3 Stop Words and Stemming	242
15.4 Using Information Gain for Feature Reduction	243
15.5 Representing Text Documents: Constructing a Vector Space Model	243
15.6 Normalising the Weights	245

15.7 Measuring the Distance Between Two Vectors	246
15.8 Measuring the Performance of a Text Classifier	247
15.9 Hypertext Categorisation	248
15.9.1 Classifying Web Pages	248
15.9.2 Hypertext Classification versus Text Classification	249
Chapter Summary	253
Self-assessment Exercises for Chapter 15	253
References	255
A. Essential Mathematics	257
A.1 Subscript Notation	257
A.1.1 Sigma Notation for Summation	258
A.1.2 Double Subscript Notation	259
A.1.3 Other Uses of Subscripts	260
A.2 Trees	260
A.2.1 Terminology	261
A.2.2 Interpretation	262
A.2.3 Subtrees	263
A.3 The Logarithm Function $\log_2 X$	264
A.3.1 The Function $-X \log_2 X$	266
A.4 Introduction to Set Theory	267
A.4.1 Subsets	269
A.4.2 Summary of Set Notation	271
B. Datasets	273
C. Sources of Further Information	293
D. Glossary and Notation	297
E. Solutions to Self-assessment Exercises	315
Index	339

Introduction to Data Mining

The Data Explosion

Modern computer systems are accumulating data at an almost unimaginable rate and from a very wide variety of sources: from point-of-sale machines in the high street to machines logging every cheque clearance, bank cash withdrawal and credit card transaction, to Earth observation satellites in space.

Some examples will serve to give an indication of the volumes of data involved:

- The current NASA Earth observation satellites generate a terabyte (i.e. 10^9 bytes) of data *every day*. This is more than the total amount of data ever transmitted by all previous observation satellites.
- The Human Genome project is storing thousands of bytes for each of several *billion* genetic bases.
- As long ago as 1990, the US Census collected over a million million bytes of data.
- Many companies maintain large Data Warehouses of customer transactions. A fairly small data warehouse might contain more than a hundred million transactions.

There are vast amounts of data recorded every day on automatic recording devices, such as credit card transaction files and web logs, as well as non-symbolic data such as CCTV recordings.

Alongside advances in storage technology, which increasingly make it possible to store such vast amounts of data at relatively low cost whether in commercial data warehouses, scientific research laboratories or elsewhere, has come

a growing realisation that such data contains buried within it knowledge that can be critical to a company's growth or decline, knowledge that could lead to important discoveries in science, knowledge that could enable us accurately to predict the weather and natural disasters, knowledge that could enable us to identify the causes of and possible cures for lethal illnesses, knowledge that could literally mean the difference between life and death. Yet the huge volumes involved mean that most of this data is merely stored—never to be examined in more than the most superficial way, if at all. It has rightly been said that the world is becoming 'data rich but knowledge poor'.

Machine learning technology, some of it very long established, has the potential to solve the problem of the tidal wave of data that is flooding around organisations, governments and individuals.

Knowledge Discovery

Knowledge Discovery has been defined as the 'non-trivial extraction of implicit, previously unknown and potentially useful information from data'. It is a process of which data mining forms just one part, albeit a central one.

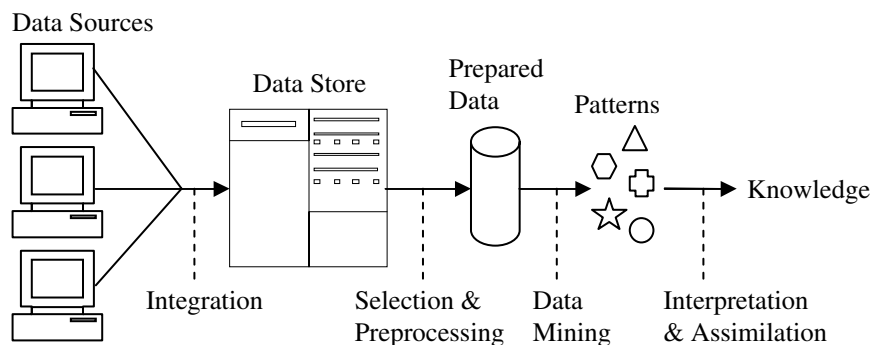


Figure 1 The Knowledge Discovery Process

Figure 1 shows a slightly idealised version of the complete knowledge discovery process.

Data comes in, possibly from many sources. It is integrated and placed in some common data store. Part of it is then taken and pre-processed into a standard format. This 'prepared data' is then passed to a data mining algorithm which produces an output in the form of rules or some other kind of 'patterns'. These are then interpreted to give—and this is the Holy Grail for knowledge discovery—new and potentially useful knowledge.

This brief description makes it clear that although the data mining algorithms, which are the principal subject of this book, are central to knowledge discovery they are not the whole story. The pre-processing of the data and the interpretation (as opposed to the blind use) of the results are both of great importance. They are skilled tasks that are far more of an art (or a skill learnt from experience) than an exact science. Although they will both be touched on in this book, the algorithms of the data mining stage of knowledge discovery will be its prime concern.

Applications of Data Mining

There is a rapidly growing body of successful applications in a wide range of areas as diverse as:

- analysis of organic compounds
- automatic abstracting
- credit card fraud detection
- electric load prediction
- financial forecasting
- medical diagnosis
- predicting share of television audiences
- product design
- real estate valuation
- targeted marketing
- thermal power plant optimisation
- toxic hazard analysis
- weather forecasting

and many more. Some examples of applications (potential or actual) are:

- a supermarket chain mines its customer transactions data to optimise targeting of high value customers
- a credit card company can use its data warehouse of customer transactions for fraud detection
- a major hotel chain can use survey databases to identify attributes of a ‘high-value’ prospect

- predicting the probability of default for consumer loan applications by improving the ability to predict bad loans
- reducing fabrication flaws in VLSI chips
- data mining systems can sift through vast quantities of data collected during the semiconductor fabrication process to identify conditions that are causing yield problems
- predicting audience share for television programmes, allowing television executives to arrange show schedules to maximise market share and increase advertising revenues
- predicting the probability that a cancer patient will respond to chemotherapy, thus reducing health-care costs without affecting quality of care.

Applications can be divided into four main types: classification, numerical prediction, association and clustering. Each of these is explained briefly below. However first we need to distinguish between two types of data.

Labelled and Unlabelled Data

In general we have a dataset of examples (called *instances*), each of which comprises the values of a number of variables, which in data mining are often called *attributes*. There are two types of data, which are treated in radically different ways.

For the first type there is a specially designated attribute and the aim is to use the data given to predict the value of that attribute for instances that have not yet been seen. Data of this kind is called *labelled*. Data mining using labelled data is known as *supervised learning*. If the designated attribute is *categorical*, i.e. it must take one of a number of distinct values such as ‘very good’, ‘good’ or ‘poor’, or (in an object recognition application) ‘car’, ‘bicycle’, ‘person’, ‘bus’ or ‘taxi’ the task is called *classification*. If the designated attribute is numerical, e.g. the expected sale price of a house or the opening price of a share on tomorrow’s stock market, the task is called *regression*.

Data that does not have any specially designated attribute is called *unlabelled*. Data mining of unlabelled data is known as *unsupervised learning*. Here the aim is simply to extract the most information we can from the data available.

Supervised Learning: Classification

Classification is one of the most common applications for data mining. It corresponds to a task that occurs frequently in everyday life. For example, a hospital may want to classify medical patients into those who are at high, medium or low risk of acquiring a certain illness, an opinion polling company may wish to classify people interviewed into those who are likely to vote for each of a number of political parties or are undecided, or we may wish to classify a student project as distinction, merit, pass or fail.

This example shows a typical situation (Figure 2). We have a dataset in the form of a table containing students' grades on five subjects (the values of attributes SoftEng, ARIN, HCI, CSA and Project) and their overall degree classifications. The row of dots indicates that a number of rows have been omitted in the interests of simplicity. We want to find some way of predicting the classification for other students given only their grade 'profiles'.

SoftEng	ARIN	HCI	CSA	Project	Class
A	B	A	B	B	Second
A	B	B	B	B	Second
B	A	A	B	A	Second
A	A	A	A	B	First
A	A	B	B	A	First
B	A	A	B	B	Second
.....
A	A	B	A	B	First

Figure 2 Degree Classification Data

There are several ways we can do this, including the following.

Nearest Neighbour Matching. This method relies on identifying (say) the five examples that are 'closest' in some sense to an unclassified one. If the five 'nearest neighbours' have grades Second, First, Second, Second and Second we might reasonably conclude that the new instance should be classified as 'Second'.

Classification Rules. We look for rules that we can use to predict the classification of an unseen instance, for example:

IF SoftEng = A AND Project = A THEN Class = First

IF SoftEng = A AND Project = B AND ARIN = B THEN Class = Second
 IF SoftEng = B THEN Class = Second

Classification Tree. One way of generating classification rules is via an intermediate tree-like structure called a *classification tree* or a *decision tree*.

Figure 3 shows a possible decision tree corresponding to the degree classification data.

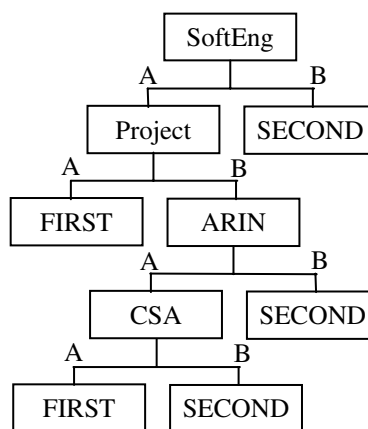


Figure 3 Decision Tree for Degree Classification Data

Supervised Learning: Numerical Prediction

Classification is one form of prediction, where the value to be predicted is a label. Numerical prediction (often called *regression*) is another. In this case we wish to predict a numerical value, such as a company's profits or a share price.

A very popular way of doing this is to use a *Neural Network* as shown in Figure 4 (often called by the simplified name *Neural Net*).

This is a complex modelling technique based on a model of a human neuron. A neural net is given a set of inputs and is used to predict one or more outputs.

Although neural networks are an important technique of data mining, they are complex enough to justify a book of their own and will not be discussed further here. There are several good textbooks on neural networks available, some of which are listed in Appendix C.




Figure 4 A Neural Network

Unsupervised Learning: Association Rules

Sometimes we wish to use a training set to find any relationship that exists amongst the values of variables, generally in the form of rules known as *association rules*. There are many possible association rules derivable from any given dataset, most of them of little or no value, so it is usual for association rules to be stated with some additional information indicating how reliable they are, for example:

IF variable_1 > 85 and switch_6 = open
THEN variable_23 < 47.5 and switch_8 = closed (probability = 0.8)

A common form of this type of application is called 'market basket analysis'. If we know the purchases made by all the customers at a store for say a week, we may be able to find relationships that will help the store market its products more effectively in the future. For example, the rule

IF cheese AND milk THEN bread (probability = 0.7)

indicates that 70% of the customers who buy cheese and milk also buy bread, so it would be sensible to move the bread closer to the cheese and milk counter, if customer convenience were the prime concern, or to separate them to encourage impulse buying of other products if profit were more important.

Unsupervised Learning: Clustering

Clustering algorithms examine data to find groups of items that are similar. For example, an insurance company might group customers according to income, age, types of policy purchased or prior claims experience. In a fault diagnosis application, electrical faults might be grouped according to the values of certain key variables (Figure 5).




Figure 5 Clustering of Data

About This Book

This book is designed to be suitable for an introductory course at either undergraduate or masters level. It can be used as a textbook for a taught unit in a degree programme on potentially any of a wide range of subjects including Computer Science, Business Studies, Marketing, Artificial Intelligence, Bioinformatics and Forensic Science. It is also suitable for use as a self-study book for those in technical or management positions who wish to gain an understanding of the subject that goes beyond the superficial. It goes well beyond the generalities of many introductory books on Data Mining but—unlike many other books—you will not need a degree and/or considerable fluency in Mathematics to understand it.

Mathematics is a language in which it is possible to express very complex and sophisticated ideas. Unfortunately it is a language in which 99% of the human race is not fluent, although many people have some basic knowledge of

it from early experiences (not always pleasant ones) at school. The author is a former Mathematician (‘recovering Mathematician’ might be a more accurate term) who now prefers to communicate in plain English wherever possible and believes that a good example is worth a hundred mathematical symbols.

Unfortunately it has not been possible to bury mathematical notation entirely. A ‘refresher’ of everything you need to know to begin studying the book is given in Appendix A. It should be quite familiar to anyone who has studied Mathematics at school level. Everything else will be explained as we come to it. If you have difficulty following the notation in some places, you can usually safely ignore it, just concentrating on the results and the detailed examples given. For those who would like to pursue the mathematical underpinnings of Data Mining in greater depth, a number of additional texts are listed in Appendix C.

No introductory book on Data Mining can take you to research level in the subject—the days for that have long passed. This book will give you a good grounding in the principal techniques without attempting to show you this year’s latest fashions, which in most cases will have been superseded by the time the book gets into your hands. Once you know the basic methods, there are many sources you can use to find the latest developments in the field. Some of these are listed in Appendix C.

The other appendices include information about the main datasets used in the examples in the book, many of which are of interest in their own right and are readily available for use in your own projects if you wish, and a glossary of the technical terms used in the book.

Self-assessment Exercises are included for each chapter to enable you to check your understanding. Specimen solutions are given in Appendix E.

Acknowledgements

I would like to thank my daughter Bryony for drawing many of the more complex diagrams and for general advice on design. I would also like to thank my wife Dawn for very valuable comments on earlier versions of the book and for preparing the index. The responsibility for any errors that may have crept into the final version remains with me.

Max Bramer
Digital Professor of Information Technology
University of Portsmouth, UK
January 2007