

Data Warehousing Mining And Olap Management Alex Berson

This book constitutes the refereed proceedings of the 15th International Conference on Data Warehousing and Knowledge Discovery, DaWaK 2013 held in Prague, Czech Republic, in August 2013. The 24 revised full papers and 8 short papers presented were carefully reviewed and selected from 89 submissions. The papers are organized in topical sections on modeling and ETL, query optimization and parallelism, spatial data warehouses and applications, text mining and OLAP, recommendation and prediction, data mining optimization and machine learning techniques, mining and processing data streams, clustering and data mining applications, social network and graph mining, and event sequence and Web mining.

This textbook covers all central activities of data warehousing and analytics, including transformation, preparation, aggregation, integration, and analysis. It discusses the full spectrum of the journey of data from operational/transactional databases, to data warehouses and data analytics; as well as the role that data warehousing plays in the data processing lifecycle. It also explains in detail how data warehouses may be used by data engines, such as BI tools and analytics algorithms to produce reports, dashboards, patterns, and other useful information and knowledge. The book is divided into six parts: “Part I – Star Schema” describes the foundation of data warehouse design. “Part II – Snowflake and Bridge Tables” then expands the concept of a simple star schema by introducing the concept of hierarchy, bridge tables, as well as the use of bridge tables in temporal data warehousing. “Part III – Advanced Dimensions” elaborates various dimension models, namely determinant dimensions, junk dimensions, dimension keys, and one-attribute dimensions, which all enrich the semantics of the star schema. “Part IV – Multi-Fact and Multi-Input” introduces multi-fact star schemas, where the star schema has multi-fact entities. A multi-fact can also be created by slicing one fact into multi-facts, which is discussed next. Eventually the creation of a star schema is introduced where an operational database is used as input to the transformation process and consists of multiple operational databases. “Part V – Data Warehousing Granularity and Evolution” first introduces the concept of aggregation levels in a star schema constellation. It then focuses on the lowest-level star schema, including how to design a star schema and why it is needed in data warehousing, before moving on to methods for adding and removing dimensions. The remaining two chapters present more advanced concepts in data warehousing granularity and introduce the concept of active data warehousing. “Part VI – OLAP, Business Intelligence, and Data Analytics” thoroughly explains OLAP – online analytical processing, and describes two important activities in the data warehousing process, namely pre-data warehousing and post-data warehousing. The final chapter focuses on data analytics, which consists of a suite of methods for data analysis suitable for data warehousing. This textbook approaches data warehousing from the case study angle. Each chapter presents one or more case studies to thoroughly explain the concepts and has different levels of difficulty, hence learning is incremental. In addition, every chapter has also a section on further readings which give pointers and references to research papers related to the chapter. All these features make the book ideally suited for either introductory courses on data warehousing and data analytics, or even for self-studies by professionals. The book is accompanied by a web page that includes all the used datasets and codes as well as slides and solutions to exercises.

The application of data warehousing and data mining techniques to computer security is an important emerging area, as information processing and internet accessibility costs decline and more and more organizations become vulnerable to cyber attacks. These security breaches include attacks on single computers, computer networks, wireless networks, databases, or authentication compromises. This book

describes data warehousing and data mining techniques that can be used to detect attacks. It is designed to be a useful handbook for practitioners and researchers in industry, and is also suitable as a text for advanced-level students in computer science.

Effective decision support systems (DSS) are quickly becoming key to businesses gaining a competitive advantage, and the effectiveness of these systems depends on the ability to construct, maintain, and extract information from data warehouses. While many still perceive data warehousing as a subdiscipline of management information systems (MIS), in fact many of its advances have and will continue to come from the computer science arena. *Intelligent Data Warehousing* presents the state of the art in data warehousing research and practice from a perspective that integrates business applications and computer science. It brings the intelligent techniques associated with artificial intelligence (AI) to the entire process of data warehousing, including data preparation, storage, and mining. Part I provides an overview of the main ideas and fundamentals of data mining, artificial intelligence, business intelligence, and data warehousing. Part II presents core materials on data warehousing, and Part III explores data analysis and knowledge discovery in the data warehousing environment, including how to perform intelligent data analysis and the discovery of influential association patterns. Bridging the gap between theoretical research and business applications, this book summarizes the main ideas behind recent research developments rather than setting forth technical details, and it presents case studies that show the how-to's of implementing these ideas. The result is a practical, first-of-its-kind book that brings together scattered research, unites MIS with computer science, and melds intelligent techniques with data warehousing.

"This book provides a comprehensive compilation of knowledge covering state-of-the-art developments and research, as well as current innovative activities in data warehousing and mining, focusing on the integration between the fields of data warehousing and data mining, with emphasis on the applicability to real world problems"--Provided by publisher.

In recent years, the science of managing and analyzing large datasets has emerged as a critical area of research. In the race to answer vital questions and make knowledgeable decisions, impressive amounts of data are now being generated at a rapid pace, increasing the opportunities and challenges associated with the ability to effectively analyze this data.

Searching for Semantics: Data Mining, Reverse Engineering Stefano Spaccapietra Fred M aryanski Swiss Federal Institute of Technology University of Connecticut Lausanne, Switzerland Storrs, CT, USA REVIEW AND FUTURE DIRECTIONS In the last few years, database semantics research has turned sharply from a highly theoretical domain to one with more focus on practical aspects. The DS-7 Working Conference held in October 1997 in Leysin, Switzerland, demonstrated the more pragmatic orientation of the current generation of leading researchers. The papers presented at the meeting emphasized the two major areas: the discovery of semantics and semantic data modeling. The work in the latter category indicates that although object-oriented database management systems have emerged as commercially viable products, many fundamental modeling issues require further investigation. Today's object-oriented systems provide the capability to describe complex objects and include techniques for mapping from a relational database to objects. However, we must further explore the expression of information regarding the dimensions of time and space. Semantic models possess the richness to describe systems containing spatial and temporal data. The challenge of incorporating these features in a manner that promotes efficient manipulation by the subject specialist still requires extensive development.

The Seventh International Symposium on Spatial and Temporal Databases (SSTD 2001), held in Redondo Beach, CA, USA, July 12-15, 2001, brought together leading researchers and developers in the area of spatial, temporal, and spatio-

temporal databases to discuss the state of the art in spatial and temporal data management and applications, and to understand the challenges and - search directions in the advancing area of data management for moving objects. The symposium served as a forum for disseminating research in spatial and temporal data management, and for maximizing the interchange of knowledge among researchers from the established spatial and temporal database communities. The exchange of research ideas and results not only contributes to the academic arena, but also benefits the user and commercial communities. SSTD 2001 was the seventh in the series of symposia that started in Santa Barbara a dozen years ago and has since been held every two years, in Zurich, Singapore, Portland (Maine), Berlin, and Hong Kong. By 1999, the series had become well established as the premier international forum devoted solely to spatial database management, and it was decided to extend the scope of the series to also cover temporal database management. This extended scope was chosen due, in part, to the increasing importance of research that considers spatial and temporal aspects jointly.

For undergraduate/graduate-level Data Mining or Data Warehousing courses in Information Systems or Operations Management Departments electives. Taking a multidisciplinary user/manager approach, this text looks at data warehousing technologies necessary to support the business processes of the twenty-first century. Using a balanced professional and conversational approach, it explores the basic concepts of data mining, warehousing, and visualization with an emphasis on both technical and managerial issues and the implication of these modern emerging technologies on those issues. Data mining and visualization exercises using an included fully-enabled, but time-limited version of Megaputer's PolyAnalyst and TextAnalyst data mining and visualization software give students hands-on experience with real-world applications.

Description: The book has been written in such a way that the concepts are explained in detail, giving adequate emphasis on examples. To make clarity on the topic, diagrams are given extensively throughout the text. The book discusses design issues for phases of mining in substantial depth. The stress is more on problem solving. Various Comprehensive coverage of various aspects of Data Mining and Warehousing concepts Strictly in accordance for the syllabus covered under B.E./B.Tech/MCA Simple language, crystal clear approach, straight forward comprehensible presentation Adopting user friendly classroom lecture style The concepts are duly supported by sever examples Syllabus coverage of three universities UPTU, RTU and RGPV Table Of Contents: Chapter 1 : Introduction To Data Mining Chapter 2 : Concept Description Chapter 3 : Association Rule Mining Chapter 4 : Classification and Predictions Chapter 5 : Cluster Analysis Chapter 6 : Introduction to Data Warehouse Chapter 7 : OLAP Technology Chapter 8 : Advance Topic On Data Mining and Warehousing

Unlike popular belief, Data Warehouse is not a single tool but a collection of software tools. A data warehouse will collect data from diverse sources into a single database. Using Business Intelligence tools, meaningful insights are drawn from this data. The best thing about "Learn Data Warehousing in 1 Day" is that it is small and can be completed in a day. With this e-book, you will be enough knowledge to contribute and participate in a Data warehouse implementation project. The book covers upcoming and promising technologies like Data Lakes, Data Mart, ETL (Extract Load Transform) amongst others. Following are detailed topics included in the book Table Of Content Chapter 1: What Is Data Warehouse? 1. What is Data Warehouse? 2. Types of Data Warehouse 3. Who needs Data warehouse? 4. Why We Need Data Warehouse? 5. Data Warehouse Tools Chapter 2: Data Warehouse Architecture 1. Characteristics of Data warehouse 2. Data Warehouse Architectures 3. Datawarehouse Components 4. Query Tools Chapter 3: ETL Process 1. What is ETL? 2. Why do you need ETL? 3. ETL Process 4. ETL tools Chapter 4: ETL Vs ELT 1. What is ETL? 2. Difference between ETL vs. ELT Chapter 5: Data Modeling 1. What is Data Modelling? 2. Types of Data Models 3. Characteristics of a physical data model Chapter 6: OLAP 1. What is Online Analytical Processing? 2. Types of OLAP systems 3. Advantages and Disadvantages of OLAP Chapter 7: Multidimensional Olap (MOLAP) 1. What is MOLAP? 2. MOLAP Architecture 3. MOLAP Tools Chapter 8: OLAP Vs OLTP 1. What is the meaning of OLAP? 2. What is the meaning of OLTP? 3. Difference between OLTP and OLAP Chapter 9: Dimensional Modeling 1. What is Dimensional Model? 2. Elements of Dimensional Data Model 3. Attributes 4. Difference between Dimension table vs. Fact table 5. Steps of Dimensional Modelling 6. Rules for Dimensional Modelling Chapter 10: Star and Snowflake Schema 1. What is Multidimensional schemas? 2. What is a Star Schema? 3. What is a Snowflake Schema? 4. Difference between Start Schema and Snowflake Chapter 11: Data Mart 1. What is Data Mart? 2. Type of Data Mart 3. Steps in Implementing a Datamart Chapter 12: Data Mart Vs Data Warehouse 1. What is Data Warehouse? 2. What is Data Mart? 3. Differences between a Data Warehouse and a Data Mart Chapter 13: Data Lake 1. What is Data Lake? 2. Data Lake Architecture 3. Key Data Lake Concepts 4. Maturity stages of Data Lake Chapter 14: Data Lake Vs Data Warehouse 1. What is Data Warehouse? 2. What is Data Lake? 3. Key Difference between the Data Lake and Data Warehouse Chapter 15: What Is Business Intelligence? 1. What is Business Intelligence 2. Why is BI important? 3. How Business Intelligence systems are implemented? 4. Four types of BI users Chapter 16: Data Mining 1. What is Data Mining? 2. Types of Data 3. Data Mining Process 4. Modelling 5. Data Mining Techniques Chapter 17: Data Warehousing Vs Data Mining 1. What is Data warehouse? 2. What Is Data Mining? 3. Difference between Data mining and Data Warehousing? Reviews planning and designing architecture and implementing the data warehouse. Includes discussions on how and why to apply IBM tools. Offers tips, tricks, and workarounds to ensure maximum performance. Companion Web site

includes technical notes, product updates, corrections, and links to relevant material and training.

Data Mining is the process of analyzing large amount of data in search of previously undiscovered business patterns.

Data Warehousing is a relational/multidimensional database that is designed for Query and Analysis rather than Transaction Processing. This book provides a systematic introduction to the principles of Data Mining and Data Warehousing. It covers the entire range of data mining algorithms (prediction, classification, and association), data mining products and applications, stages.

Data Mining: Concepts and Techniques provides the concepts and techniques in processing gathered data or information, which will be used in various applications. Specifically, it explains data mining and the tools used in discovering knowledge from the collected data. This book is referred as the knowledge discovery from data (KDD). It focuses on the feasibility, usefulness, effectiveness, and scalability of techniques of large data sets. After describing data mining, this edition explains the methods of knowing, preprocessing, processing, and warehousing data. It then presents information about data warehouses, online analytical processing (OLAP), and data cube technology. Then, the methods involved in mining frequent patterns, associations, and correlations for large data sets are described. The book details the methods for data classification and introduces the concepts and methods for data clustering. The remaining chapters discuss the outlier detection and the trends, applications, and research frontiers in data mining. This book is intended for Computer Science students, application developers, business professionals, and researchers who seek information on data mining. Presents dozens of algorithms and implementation examples, all in pseudo-code and suitable for use in real-world, large-scale data mining projects Addresses advanced topics such as mining object-relational databases, spatial databases, multimedia databases, time-series databases, text databases, the World Wide Web, and applications in several fields Provides a comprehensive, practical look at the concepts and techniques you need to get the most out of your data

"Data Warehousing" is the nuts-and-bolts guide to designing a data management system using data warehousing, data mining, and online analytical processing (OLAP) and how successfully integrating these three technologies can give business a competitive edge.

The field of data mining provides techniques for automated discovery of valuable information from the accumulated data of computerized operations of enterprises. This book offers a clear and comprehensive introduction to both data mining theory and practice. It is written primarily as a textbook for the students of computer science, management, computer applications, and information technology. The book ensures that the students learn the major data mining techniques even if they do not have a strong mathematical background. The techniques include data pre-processing, association rule mining, supervised classification,

cluster analysis, web data mining, search engine query mining, data warehousing and OLAP. To enhance the understanding of the concepts introduced, and to show how the techniques described in the book are used in practice, each chapter is followed by one or two case studies that have been published in scholarly journals. Most case studies deal with real business problems (for example, marketing, e-commerce, CRM). Studying the case studies provides the reader with a greater insight into the data mining techniques. The book also provides many examples, review questions, multiple choice questions, chapter-end exercises and a good list of references and Web resources especially those which are easy to understand and useful for students. A number of class projects have also been included.

This Book Is Mainly Intended For It Students And Professionals To Learn Or Implement Data Warehousing Technologies. It Experiences The Real-Time Environment And Promotes Planning, Managing, Designing, Implementing, Supporting, Maintaining And Analyzing Data Warehouse In Organizations And It Also Provides Various Mining Techniques As Well As Issues In Practical Use Of Data Mining Tools. The Book Is Designed For The Target Audience Such As Specialists, Trainers And It Users. It Does Not Assume Any Special Knowledge As Background. Understanding Of Computer Use, Databases And Statistics Will Be Helpful. There are more than one billion documents on the Web, with the count continually rising at a pace of over one million new documents per day. As information increases, the motivation and interest in data warehousing and mining research and practice remains high in organizational interest. The Encyclopedia of Data Warehousing and Mining, Second Edition, offers thorough exposure to the issues of importance in the rapidly changing field of data warehousing and mining. This essential reference source informs decision makers, problem solvers, and data mining specialists in business, academia, government, and other settings with over 300 entries on theories, methodologies, functionalities, and applications.

Data warehousing is one of the hottest business topics, and there's more to understanding data warehousing technologies than you might think. Find out the basics of data warehousing and how it facilitates data mining and business intelligence with Data Warehousing For Dummies, 2nd Edition. Data is probably your company's most important asset, so your data warehouse should serve your needs. The fully updated Second Edition of Data Warehousing For Dummies helps you understand, develop, implement, and use data warehouses, and offers a sneak peek into their future. You'll learn to: Analyze top-down and bottom-up data warehouse designs Understand the structure and technologies of data warehouses, operational data stores, and data marts Choose your project team and apply best development practices to your data warehousing projects Implement a data warehouse, step by step, and involve end-users in the process Review and upgrade existing data storage to make it serve your needs Comprehend OLAP, column-wise databases, hardware assisted databases, and middleware Use data mining intelligently and find what you need Make informed choices about consultants and data warehousing products Data Warehousing For Dummies, 2nd Edition also shows you how to involve users in the testing process and gain valuable feedback, what it takes to successfully manage a data warehouse project, and how to tell if your project is on track. You'll find it's the most useful source of data on the topic!

CUTTING-EDGE CONTENT AND GUIDANCE FROM A DATA WAREHOUSING EXPERT—NOW EXPANDED TO REFLECT FIELD TRENDS Data warehousing has revolutionized the way businesses in a wide variety of industries perform analysis and make strategic decisions. Since the first edition of *Data Warehousing Fundamentals*, numerous enterprises have implemented data warehouse systems and reaped enormous benefits. Many more are in the process of doing so. Now, this new, revised edition covers the essential fundamentals of data warehousing and business intelligence as well as significant recent trends in the field. The author provides an enhanced, comprehensive overview of data warehousing together with in-depth explanations of critical issues in planning, design, deployment, and ongoing maintenance. IT professionals eager to get into the field will gain a clear understanding of techniques for data extraction from source systems, data cleansing, data transformations, data warehouse architecture and infrastructure, and the various methods for information delivery. This practical Second Edition highlights the areas of data warehousing and business intelligence where high-impact technological progress has been made. Discussions on developments include data marts, real-time information delivery, data visualization, requirements gathering methods, multi-tier architecture, OLAP applications, Web clickstream analysis, data warehouse appliances, and data mining techniques. The book also contains review questions and exercises for each chapter, appropriate for self-study or classroom work, industry examples of real-world situations, and several appendices with valuable information. Specifically written for professionals responsible for designing, implementing, or maintaining data warehousing systems, *Data Warehousing Fundamentals* presents agile, thorough, and systematic development principles for the IT professional and anyone working or researching in information management. Geared to IT professionals eager to get into the all-important field of data warehousing, this book explores all topics needed by those who design and implement data warehouses. Readers will learn about planning requirements, architecture, infrastructure, data preparation, information delivery, implementation, and maintenance. They'll also find a wealth of industry examples garnered from the author's 25 years of experience in designing and implementing databases and data warehouse applications for major corporations. Market: IT Professionals, Consultants.

Provides developments and research, as well as current innovative activities in data warehousing and mining, focusing on the intersection of data warehousing and business intelligence.

The Skills You Need to Develop OLAP Solutions with SQL Server 2000 This one-of-a-kind book teaches you everything you need to know to use Microsoft's Analysis Services software to build, implement, and manage effective OLAP solutions. Expert advice and in-depth explanations combine to help you and your company take full advantage of the affordable power of SQL Server's built-in OLAP functionality. Coverage Includes: Analyzing large volumes of data effectively with Analysis Services Architecting and designing data analysis applications Querying OLAP data using MDX Programming applications using ADO/MD Managing Analysis Services servers with DSO Building data mining solutions with Analysis Services Using English Query for natural language querying of OLAP data Choosing appropriate client tools for exploring OLAP data Using the PivotTable Service for client-side data analysis Note: CD-ROM/DVD and other supplementary materials are not included as part of eBook file.

"This book provides insight into the latest findings concerning data warehousing, data mining, and their applications in everyday human activities"--Provided by publisher.

Data warehouses and online analytical processing (OLAP) are emerging key technologies for enterprise decision support systems. They provide sophisticated technologies from data integration, data collection and retrieval, query optimization, and data analysis to advanced user interfaces. New research and technological achievements in the area of data warehousing are implemented in commercial database management systems, and organizations are developing data warehouse systems into their information system infrastructures. *Data Warehouses and OLAP: Concepts, Architectures and Solutions* covers a wide range of technical, technological, and research issues. It provides theoretical frameworks, presents challenges and their possible solutions, and examines the latest empirical research findings in the area. It is a resource of possible solutions and technologies that can be applied when designing, implementing, and deploying a data warehouse, and assists in the dissemination of knowledge in this field.

Data Warehousing and Data Mining is presented in a question-and-answer format following the examination pattern and covers all key topics in the syllabus. The book is designed to make learning fast and effective and is precise, up-to-date and will help students excel in their examinations. The book is part of the Express Learning is a series of books designed as quick reference guides to important undergraduate courses. The organized and accessible format of these books allows students to learn important concepts in an easy-to-understand, question-and-answer format. These portable learning tools have been designed as one-stop references for students to understand and master the subjects by themselves.

This book constitutes the refereed proceedings of the 14th International Conference on Data Warehousing and Knowledge Discovery, DaWaK 2012 held in Vienna, Austria, in September 2012. The 36 revised full papers presented were carefully reviewed and selected from 99 submissions. The papers are organized in topical sections on data warehouse design methodologies, ETL methodologies and tools, multidimensional data processing and management, data warehouse and OLAP extensions, data warehouse performance and optimization, data mining and knowledge discovery techniques, data mining and knowledge discovery applications, pattern mining, data stream mining, data warehouse confidentiality and security, and distributed paradigms and algorithms.

This book constitutes the refereed proceedings of the Second International Conference on Advances in Communication, Network, and Computing, CNC 2011, held in Bangalore, India, in March 2011. The 41 revised full papers, presented together with 50 short papers and 39 poster papers, were carefully reviewed and selected for inclusion in the book. The papers feature current research in the field of Information Technology, Networks, Computational Engineering, Computer and Telecommunication Technology, ranging from theoretical and methodological issues to advanced applications.

This book constitutes the refereed proceedings of the 10th International Conference on Data Warehousing and Knowledge Discovery, DaWaK 2008, held in Turin, Italy, in September 2008. The 40 revised full papers presented were carefully reviewed and

selected from 143 submissions. The papers are organized in topical sections on conceptual design and modeling, olap and cube processing, distributed data warehouse, data privacy in data warehouse, data warehouse and data mining, clustering, mining data streams, classification, text mining and taxonomy, machine learning techniques, and data mining applications.

Open access to information of geographic places and spatial relationships provides an essential part of the analytical processing of spatial data. Access to connected geospatial programs allows for improvement in teaching and understanding science, technology, engineering, and mathematics. Emerging Trends in Open Source Geographic Information Systems provides emerging research on the applications of free and open software in geographic information systems in various fields of study. While highlighting topics such as data warehousing, hydrological modeling, and software packages, this publication explores the assessment and techniques of open software functionality and interfaces. This book is an important resource for professionals, researchers, academicians, and students seeking current research on the different types and uses of data and data analysis in geographic information systems.

This book constitutes the refereed proceedings of the 16th International Conference on Data Warehousing and Knowledge Discovery, DaWaK 2014 held in Munich, Germany, September 2014, in conjunction with DEXA 2014. The 34 revised full papers and 8 short papers presented were carefully reviewed and selected from 109 submissions. The papers are organized in topical sections on modeling and ETL; ontology-based data warehouses; advanced data warehouses and OLAP; uncertainty; preferences and recommendation; query performance and HPC; cube & OLAP; optimization; classification; social networks and recommendation systems; knowledge data discovery; industrial applications; mining and processing data stream; mining and similarity.

This book constitutes the refereed proceedings of the 13th International Conference on Data Warehousing and Knowledge Discovery, DaWaK 2011 held in Toulouse, France in August/September 2011. The 37 revised full papers presented were carefully reviewed and selected from 119 submissions. The papers are organized in topical sections on physical and conceptual data warehouse models, data warehousing design methodologies and tools, data warehouse performance and optimization, pattern mining, matrix-based mining techniques and stream, sensor and time-series mining.

Written in lucid language, this valuable textbook brings together fundamental concepts of data mining and data warehousing in a single volume. Important topics including information theory, decision tree, Naive Bayes classifier, distance metrics, partitioning clustering, associate mining, data marts and operational data store are discussed comprehensively. The textbook is written to cater to the needs of undergraduate students of computer science, engineering and information technology for a course on data mining and data warehousing. The text simplifies the understanding of the concepts through exercises and practical examples. Chapters such as classification, associate mining and cluster analysis are discussed in detail with their practical implementation using Weka and R language data mining tools. Advanced topics including big data analytics, relational data models and NoSQL are discussed in detail. Pedagogical features including unsolved problems and multiple-choice questions are interspersed

throughout the book for better understanding.

This book constitutes the refereed proceedings of the 11th International Conference on Data Warehousing and Knowledge Discovery, DaWak 2009 held in Linz, Austria in August/September 2009. The 36 revised full papers presented were carefully reviewed and selected from 124 submissions. The papers are organized in topical sections on data warehouse modeling, data streams, physical design, pattern mining, data cubes, data mining applications, analytics, data mining, clustering, spatio-temporal mining, rule mining, and OLAP recommendation.

Within the last few years Data Warehousing and Knowledge Discovery technology has established itself as a key technology for enterprises that wish to improve the quality of the results obtained from data analysis, decision support, and the automatic extraction of knowledge from data. The Fourth International Conference on Data Warehousing and Knowledge Discovery (DaWaK 2002) continues a series of successful conferences dedicated to this topic. Its main objective is to bring together researchers and practitioners to discuss research issues and experience in developing and deploying data warehousing and knowledge discovery systems, applications, and solutions. The conference focuses on the logical and physical design of data warehousing and knowledge discovery systems. The scope of the papers covers the most recent and relevant topics in the areas of association rules, clustering, Web mining, security, data mining techniques, data cleansing, applications, data warehouse design and maintenance, and OLAP. These proceedings contain the technical papers selected for presentation at the conference. We received more than 100 papers from over 20 countries, and the program committee finally selected 32 papers. The conference program included one invited talk: "Text Mining Applications of a Shallow Parser" by Walter Daelemans, University of Antwerp, Belgium. We would like to thank the DEXA 2002 Workshop General Chair (Roland Wagner) and the organizing committee of the 13 International Conference on Database and Expert Systems Applications (DEXA 2002) for their support and their cooperation.

[Copyright: a745aa0a4a4ce7b73ea70754bff2dbeb](#)