Relational Data Is Based On Which Three Mathematical Concepts

BROCK HARVEY

Granular-Relational Data Mining Springer Science & Business Media
Ranking queries are widely used in data exploration, data analysis and decision making scenarios. While most of the currently proposed ranking techniques focus on deterministic data, several emerging applications involve data that are imprecise or uncertain. Ranking uncertain data raises new challenges in query semantics and processing, making conventional methods inapplicable. Furthermore, the interplay between ranking and uncertainty models introduces new dimensions for ordering query results that do not exist in the traditional settings. This lecture describes new formulations and processing techniques for ranking queries on uncertain data. The formulations are based on marriage of traditional ranking semantics with possible worlds semantics under widely-adopted uncertainty models. In particular, we focus on discussing the impact of tuple-level and attribute-level uncertainty on the semantics and processing techniques of ranking queries. Under the tuple-level uncertainty model, we describe new processing techniques leveraging the capabilities of relational database systems to recognize and handle data uncertainty in score-based ranking. Under the attribute-level uncertainty model, we describe new probabilistic ranking models and a set of query evaluation algorithms, including sampling-based techniques. We also discuss supporting rank join queries on uncertain data, and we show how to extend current rank join methods to handle uncertainty in scoring attributes. Table of Contents: Introduction / Uncertainty Models / Query Semantics / Methodologies / Uncertain Rank Join / Conclusion
Datalog and Logic Databases Springer Nature
This book provides comprehensive coverage of fundamentals of database management system. It contains a detailed description on Relational Database Management System Concepts. There are a variety of solved examples and review questions with solutions. This book is for those who require a better understanding of relational data modeling, its purpose, its nature, and the standards used in creating relational data model.
Multilevel Security for Relational Databases Addison Wesley Publishing Company
Introduced forty years ago, relational databases proved unusually succe- ful and durable. However, relational database systems were not designed for modern applications and computers. As a result, specialized database systems now proliferate trying to capture various pieces of the database market. Database research is pulled into di?erent directions, and speci- ized database conferences are created. Yet the current chaos in databases is likely only temporary because every technology, including databases, becomes standardized over time. The history of databases shows periods of chaos followed by periods of dominant technologies. For example, in the early days of computing, users stored their data in text ?les in any format and organization they wanted. These early days were followed by information retrieval systems, which required some structure for text documents, such as a title, authors, and a publisher. The information retrieval systems were followed by database systems, which added even more structure to the data and made querying easier. In the late 1990s, the emergence of the Internet brought a period of relative chaos and interest in unstructured and “semistructured data” as it wasenvisionedthateverywebpagewouldbelikeapageinabook.However, with the growing maturity of the Internet, the interest in structured data was regained because the most popular websites are, in fact, based on databases. The question is not whether future data stores need structure but what structure they need.
Integrating Relational Databases with the Semantic Web Springer Science & Business Media
Learn the best way of writing code to run inside a relational database. This book shows how a holistic and set-oriented approach to database programming can far exceed the performance of the row-by-row model that is too often used by developers who haven’t been shown a better way. Two styles of programming are encountered in the database world. Classical programming as taught in many universities leads to an atomic, row-oriented, and procedural style inspired by the structured models of programming. In short, many application developers write in the relational database exactly like in the user interface. The other style of programming is holistic, data set oriented, and coded mainly in SQL. This is the style of the database developer. The set based and holistic style of development is not promoted enough in universities, and many application developers are not fully aware of it. There are many performance issues all over the world in relational databases due to the use of the atomic and inappropriate style of programming. This book compares the two styles, and promotes the holistic style of development as the most suitable one. Examples are given to demonstrate the superiority of a set-based and holistic approach. Compares the two styles of development Shows the performance advantages of set-based development Solves example problems using both approaches Who This Book Is For Two Styles of Database Development is aimed at application developers willing to adapt their programming styles in return for better-performing applications. It’s for students and new developers wanting to position themselves as having database expertise and build a reputation for developing highly-performant database applications.
Relational Database Design and Implementation Morgan & Claypool Publishers
Relational Database Design and Implementation: Clearly Explained, Fourth Edition, provides the conceptual and practical information necessary to develop a database design and management scheme that ensures data accuracy and user satisfaction while optimizing performance. Database systems underlie the large majority of business information systems. Most of those in use today are based on the relational data model, a way of representing data and data relationships using only two-dimensional tables. This book covers relational database theory as well as providing a solid introduction to SQL, the international standard for the relational database data manipulation language. The book begins by reviewing basic concepts of databases and database design, then turns to creating, populating, and retrieving data using SQL. Topics such as the relational data model, normalization, data entities, and Codd's Rules (and why they are important) are covered clearly and concisely. In addition, the book looks at the impact of big data on relational databases and the option of using NoSQL databases for that purpose. Features updated and expanded coverage of SQL and new material on big data, cloud computing, and object-relational databases Presents design approaches that ensure data accuracy and consistency and help boost performance Includes three case studies, each illustrating a different database design challenge Reviews the basic concepts of databases and database design, then turns to creating, populating, and retrieving data using SQL
Introduction to Databases Apress
Logical Data Modeling offers business managers, analysts, and students a clear, basic systematic guide to defining business information structures in relational database terms. The approach, based on Clive Finkelstein’s business-side Information Engineering, is hands-on, practical, and explicit in terminology and reasoning. Filled with illustrations, examples, and exercises, Logical Data Modeling makes its subject accessible to readers with only a limited knowledge of database systems. The book covers all essential topics thoroughly but succinctly: entities, associations, attributes, keys and inheritance, valid and invalid structures, and normalization. It also emphasizes communication with business and database specialists, documentation, and the use of Visible Systems' Visible Advantage enterprise modeling tool. The application of design patterns to logical data modeling provides practitioners with a practical tool for fast development. At the end, a chapter covers the issues that arise when the logical data model is translated into the design for a physical database.
Relational Database Programming CRC Press
Design and Use of Relational Databases in Chemistry helps programmers and users improve their ability to search and manipulate chemical structures and information. It illustrates how

the organizational, data integrity, and extensibility properties of relational databases are best utilized when working with chemical information. The author facilitates an understanding of existing relational database schemas and shows how to design new schemas that contain tables of data and chemical structures. By using database extension "cartridges," he provides methods to properly store and search chemical structures. He explains how to download and install a fully functioning database using free, open-source chemical extension cartridges. The author also discusses how to access a database on a computer network using both new and existing applications. Through examples of good database design, this book shows you that relational databases are the best way to store, search, and operate on chemical information. Features, Introduces the concepts of relational database management systems and the structured query language, Discusses how relational data tables and new data types help store and use chemical information, Shows how molecular structures can become a new data type in a database, Describes how client programs, including web-based applications, can effectively use relational databases, Explains how a fully functioning chemical relational database system can be built, Includes many practical examples of experimental, theoretical, and structural chemical data relations, Provides a website that offers an implementation of every function from the book, a database of structures and data, and examples used in the book, allowing you to experiment with various search and display options Book jacket.

Similarity Joins in Relational Database Systems Springer Science & Business Media

A culmination of the authors' years of extensive research on this topic, Relational Data Clustering: Models, Algorithms, and Applications addresses the fundamentals and applications of relational data clustering. It describes theoretic models and algorithms and, through examples, shows how to apply these models and algorithms to solve real-world problems. After defining the field, the book introduces different types of model formulations for relational data clustering, presents various algorithms for the corresponding models, and demonstrates applications of the models and algorithms through extensive experimental results. The authors cover six topics of relational data clustering: Clustering on bi-type heterogeneous relational data Multi-type heterogeneous relational data Homogeneous relational data clustering Clustering on the most general case of relational data Individual relational clustering framework Recent research on evolutionary clustering This book focuses on both practical algorithm derivation and theoretical framework construction for relational data clustering. It provides a complete, self-contained introduction to advances in the field.

Improving data quality in relational databases Apress

Addressing important extensions of the relational database model, including deductive, temporal, and object-oriented databases, this book provides an overview of database modeling with the Entity-Relationship (ER) model and the relational model. The book focuses on the primary achievements in relational database theory, including query languages, integrity constraints, database design, computable queries, and concurrency control. This reference will shed light on the ideas underlying relational database systems and the problems that confront database designers and researchers.

SQL Server Big Data Clusters CRC Press

Relational data exchange is the problem of translating relational data according to a given specification. It is one of the many tasks that arise in information integration. A fundamental issue is how to answer queries that are posed against the result of the data exchange so that the answers are semantically consistent with the source data. For monotonic queries, the certain answers semantics by Fagin, Kolaitis, Miller, and Popa (2003) yields good answers. For many non-monotonic queries, however, this semantics was shown to yield counter-intuitive answers. This dissertation deals with the problem of computing the certain answers to monotonic queries on the one hand. On the other hand, it presents and compares semantics for answering non-monotonic queries, and investigates how hard it is to evaluate non-monotonic queries under these semantics.

Keyword Search in Databases "O'Reilly Media, Inc."

Create database designs that scale, meet business requirements, and inherently work toward keeping your data structured and usable in the face of changing business models and software systems. This book is about database design theory. Design theory is the scientific foundation for database design, just as the relational model is the scientific foundation for database technology in general. Databases lie at the heart of so much of what we do in the computing world that negative impacts of poor design can be extraordinarily widespread. This second edition includes greatly expanded coverage of exotic and little understood normal forms such as: essential tuple normal form (ETNF), redundancy free normal form (RFNF), superkey normal form (SKNF), sixth normal form (6NF), and domain key normal form (DKNF). Also included are new appendixes, including one that provides an in-depth look into the crucial notion of data consistency. Sequencing of topics has been improved, and many explanations and examples have been rewritten and clarified based upon the author's teaching of the content in instructor-led courses. This book aims to be different from other books on design by bridging the gap between the theory of design and the practice of design. The book explains theory in a way that practitioners should be able to understand, and it explains why that theory is of considerable practical importance. Reading this book provides you with an important theoretical grounding on which to do the practical work of database design. Reading the book also helps you in going to and understanding the more academic texts as you build your base of knowledge and expertise. Anyone with a professional interest in database design can benefit from using this book as a stepping-stone toward a more rigorous design approach and more lasting database models. What You Will Learn Understand what design theory is and is not Be aware of the two different goals of normalization Know which normal forms are truly significant Apply design theory in practice Be familiar with techniques for dealing with redundancy Understand what consistency is and why it is crucially important Who This Book Is For Those having a professional interest in database design, including data and database administrators; educators and students specializing in database matters; information modelers and database designers; DBMS designers, implementers, and other database vendor personnel; and database consultants. The book is product independent.

Six-Step Relational Database Design Springer Science & Business Media

The traditional vertical decomposition methods in relational database normalization fail to prevent common data anomalies. Although a database may be highly normalized, the quality of the data stored in this database may still deteriorate because of potential data anomalies. In this paper, we first discuss why practitioners need to further improve their databases after they apply the traditional normalization methods, because of the existence of functional entanglement, a phenomenon we defined. We outline two methods for identifying functional entanglements in a normalized database as the first step toward data quality improvement. We then analyze several practical methods for preventing common data anomalies by eliminating and restricting the effects of functional entanglements. The goal of this paper is to reveal shortcomings of the traditional database normalization methods with respect to the prevention of common data anomalies, and offer practitioners useful techniques for improving data quality.

Linked Data Management IOS Press

This is a book on database management that is based on an earlier book by the same authors,Foundation for Future Database Systems: The Third Manifesto. It can be seen as an abstract blueprint for the design of a DBMS and the language interface to such a DBMS. In particular, it serves as a basis for a model of type inheritance. This book is essential reading for database professionals.

Beginning Relational Data Modeling Springer

R for Data Science"O'Reilly Media, Inc."

Theory and Practice of Relational Databases Fidel A Captain

Ranking queries are widely used in data exploration, data analysis and decision making scenarios. While most of the currently proposed ranking techniques focus on deterministic data, several emerging applications involve data that are imprecise or uncertain. Ranking uncertain data raises new challenges in query semantics and processing, making conventional methods inapplicable. Furthermore, the interplay between ranking and uncertainty models introduces new dimensions for ordering query results that do not exist in the traditional settings. This lecture describes new formulations and processing techniques for ranking queries on uncertain data. The formulations are based on marriage of traditional ranking semantics with possible worlds semantics under widely-adopted uncertainty models. In particular, we focus on discussing the impact of tuple-level and attribute-level uncertainty on the semantics and processing techniques of ranking queries. Under the tuple-level uncertainty model, we describe new processing techniques leveraging the capabilities of relational database systems to recognize and handle data uncertainty in score-based ranking. Under the attribute-level uncertainty model, we describe new probabilistic ranking models and a set of query evaluation algorithms, including sampling-based techniques. We also discuss supporting rank join queries on uncertain data, and we show how to extend current rank join methods to handle uncertainty in scoring attributes. Table of Contents: Introduction / Uncertainty Models / Query Semantics / Methodologies / Uncertain Rank Join / Conclusion

Fuzzy Preference Queries to Relational Databases Springer Nature

This book introduces the fundamental concepts necessary for designing, using, and implementing database systems and database applications. Our presentation stresses the fundamentals of database modeling and design, the languages and models provided by the database management systems, and database system implementation techniques. The book is meant to be used as a textbook for a one- or two-semester course in database systems at the junior, senior, or graduate level, and as a reference book. Our goal is to provide an in-depth and up-to-date presentation of the most important aspects of database systems and applications, and related technologies. We assume that readers are familiar with elementary programming and data structuring concepts and those they have had some exposure to the basics of computer organization.

Data Analysis Using SQL and Excel CRC Press

As the first book devoted to relational data mining, this coherently written multi-author monograph provides a thorough introduction and systematic overview of the area. The first part introduces the reader to the basics and principles of classical knowledge discovery in databases and inductive logic programming; subsequent chapters by leading experts assess the techniques in relational data mining in a principled and comprehensive way; finally, three chapters deal with advanced applications in various fields and refer the reader to resources for relational data mining. This book will become a valuable source of reference for R&D professionals active in relational data mining. Students as well as IT professionals and ambitioned practitioners interested in learning about relational data mining will appreciate the book as a useful text and gentle introduction to this exciting new field.

Design and Use of Relational Databases in Chemistry Springer Science & Business Media

Contents Should we tell you the whole story? Of course, there is an inevitable tension in trying to work like this. For example, in Chapter 16 we talk about referential integrity. There are - sentially six different flavors of referential integrity but Access only s- ports four of them (they are the most important ones however, so you aren't missing out on too much). The problem is this. Should we tell you about the other two? If we do, as an Access user you have every right to be annoyed that we are telling you about a feature you can't use. On the other hand, the six different types that we describe are part of the re- tional world and this book is about that world – we are not trying to teach you how to use Access, we are simply using Access to illustrate the relational model. Ultimately we decided to risk your ire and to describe all of the features of the relational model as we see it, even if Access doesn't support all of them. One advantage of this approach is that if you need to use a different database engine you will almost certainly find the extra information useful. Incidentally, this is not meant to imply that Access is somehow lacking as a relational database engine. The reason we chose it for the first book is that it is such a good example of a relational database tool.

Relational Data Clustering John Wiley & Sons

Get a head-start on learning one of SQL Server 2019's latest and most impactful features—Big Data Clusters—that combines large volumes of non-relational data for analysis along with data stored relationally inside a SQL Server database. This book provides a first look at Big Data Clusters based upon SQL Server 2019 Release Candidate 1. Start now and get a jump on your competition in learning this important new feature. Big Data Clusters is a feature set covering data virtualization, distributed computing, and relational databases and provides a complete AI platform across the entire cluster environment. This book shows you how to deploy, manage, and use Big Data Clusters. For example, you will learn how to combine data stored on the HDFS file system together with data stored inside the SQL Server instances that make up the Big Data Cluster. Filled with clear examples and use cases, this book provides everything necessary to get started working with Big Data Clusters in SQL Server 2019 using Release Candidate 1. You will learn about the architectural foundations that are made up from Kubernetes, Spark, HDFS, and SQL Server on Linux. You then are shown how to configure and deploy Big Data Clusters in on-premises environments or in the cloud. Next, you are taught about querying. You will learn to write queries in Transact-SQL—taking advantage of skills you have honed for years—and with those queries you will be able to examine and analyze data from a wide variety of sources such as Apache Spark. Through the theoretical foundation provided in this book and easy-to-follow example scripts and notebooks, you will be ready to use and unveil the full potential of SQL Server 2019: combining different types of data spread across widely disparate sources into a single view that is useful for business intelligence and machine learning analysis. What You Will LearnInstall, manage, and troubleshoot Big Data Clusters in cloud or on-premise environments Analyze large volumes of data directly from SQL Server and/or Apache Spark Manage data stored in HDFS from SQL Server as if it were relational data Implement advanced analytics solutions through machine learning and AI Expose different data sources as a single logical source using data virtualization Who This Book Is For For data engineers, data scientists, data architects, and database administrators who want to employ data virtualization and big data analytics in their environment

R for Data Science

Since databases are the primary repositories of information for today's organizations and governments, database security has become critically important. Introducing the concept of multilevel security in relational databases, this book provides a comparative study of the various models that support multilevel security policies in the relational database—illustrating the strengths and weaknesses of each model. Multilevel Security for Relational Databases covers multilevel database security concepts along

with many other multilevel database security models and techniques. It presents a prototype that readers can implement as a tool for conducting performance evaluations to compare multilevel secure database models. The book supplies a complete view of an encryption-based multilevel security database model that integrates multilevel security for the relational database with a system that encrypts each record with an encryption key according to its security class level. This model will help you utilize an encryption system as a second security layer over the multilevel security layer for the database, reduce the multilevel database size, and improve the response time of data retrieval from the multilevel database. Considering instance-based multilevel database security, the book covers relational database access controls and examines concurrency control in multilevel database security systems. It includes database encryption algorithms, simulation programs, and Visual studio and Microsoft SQL Server code.