Data Matching Concepts And Techniques For Record Linkage Entity Resolution And Duplicate Detection Data Centric Systems And Applications 2012 Edition By Christen Peter Published By Springer 2012

Yeah, reviewing a books **Data Matching Concepts And Techniques For Record Linkage Entity Resolution And Duplicate Detection Data Centric Systems And Applications 2012 Edition By Christen Peter Published By Springer 2012** could ensue your close contacts listings. This is just one of the solutions for you to be successful. As understood, skill does not recommend that you have fabulous points.

Comprehending as well as contract even more than extra will have enough money each success. next-door to, the notice as without difficulty as insight of this Data Matching Concepts And Techniques For Record Linkage Entity Resolution And Duplicate Detection Data Centric Systems And Applications 2012 Edition By Christen Peter Published By Springer 2012 can be taken as competently as picked to act.

Big Data and Social Science - Ian Foster 2016-08-10
Both Traditional Students and Working Professionals Acquire the Skills to Analyze Social Problems. Big Data and Social Science: A Practical Guide to Methods and Tools shows how to apply data science to real-world problems in both research and the practice. The book provides practical guidance on combining methods and tools from computer science, statistics, and social science. This concrete approach is illustrated throughout using an important national problem, the quantitative study of innovation. The text draws on the expertise of prominent leaders in statistics, the social sciences, data science, and computer science to teach students how to use modern social science research principles as well as the best analytical and computational tools. It uses a real-world challenge to introduce how these tools are used to identify and capture appropriate data, apply data science models and tools to that data, and recognize and respond to data errors and limitations. For more information, including sample chapters and news, please visit the author's website.

Model Rules of Professional Conduct - American Bar Association. House of Delegates 2007
The Model Rules of Professional Conduct provides an up-to-date resource for information on legal ethics. Federal, state and local courts in all jurisdictions look to the Rules for guidance in solving lawyer malpractice cases, disciplinary actions, disqualification issues, sanctions questions and much more. In this volume, black-letter Rules of Professional Conduct are followed by numbered Comments that explain each Rule's purpose and provide suggestions for its

practical application. The Rules will help you identify proper conduct in a variety of given situations, review those instances where discretionary action is possible, and define the nature of the relationship between you and your clients, colleagues and the courts.

The Practitioner's Guide to Graph Data - Denise Gosnell 2020-03-20

Graph data closes the gap between the way humans and computers view the world. While computers rely on static rows and columns of data, people navigate and reason about life through relationships. This practical guide demonstrates how graph data brings these two approaches together. By working with concepts from graph theory, database schema, distributed systems, and data analysis, you'll arrive at a unique intersection known as graph thinking. Authors Denise Koessler Gosnell and Matthias Broecheler show data engineers, data scientists, and data analysts how to solve complex problems with graph databases. You'll explore templates for building with graph technology, along with examples that demonstrate how teams think about graph data within an application. Build an example application architecture with relational and graph technologies Use graph technology to build a Customer 360 application, the most popular graph data pattern today Dive into hierarchical data and troubleshoot a new paradigm that comes from working with graph data Find paths in graph data and learn why your trust in different paths motivates and informs your preferences Use collaborative filtering to design a Netflix-inspired recommendation system

Handbook on Impact Evaluation - Shahidur R. Khandker 2009-10-13

Public programs are designed to reach certain goals and beneficiaries. Methods

to understand whether such programs actually work, as well as the level and nature of impacts on intended beneficiaries, are main themes of this book.

IBM Classification Module: Make It Work for You - Wei-Dong Zhu 2009-11-03

IBM® Classification Module (Classification Module) Version 8.6 is an advanced enterprise software platform tool designed to allow organizations to automate the classification of unstructured content. By deploying the module in various areas of a business, organizations can reduce or avoid manual processes associated with subjective decision making around unstructured content. Organizations can also streamline the ingestion of that content into their business systems in order to use the information within the business systems more effectively. At the same time, the organizations can safely remove irrelevant or obsolete information and therefore utilize the storage infrastructure more efficiently. By reducing the human element in this process, Classification Module ensures accuracy and consistency and enables auditing while simultaneously driving down labor costs. This IBM Redbooks® publication explains what Classification Module does, the key concepts to understand when working with Classification Module, and its integration with other products and systems. With this book, we show you how Classification Module helps your organization to automate the classification of large volumes of unstructured content in a consistent and accurate manner. The topics that are covered include building, training, and fine-tuning the knowledge base, creating decision plans, working with Classification Workbench, and step-by-step integration with

other products and solutions. This book is intended to educate both technical specialists and nontechnical personnel in how to make Classification Module work for your organizations.

Data-Driven Policy Impact Evaluation - Nuno Crato 2018-10-02

In the light of better and more detailed administrative databases, this open access book provides statistical tools for evaluating the effects of public policies advocated by governments and public institutions. Experts from academia, national statistics offices and various research centers present modern econometric methods for an efficient data-driven policy evaluation and monitoring, assess the causal effects of policy measures and report on best practices of successful data management and usage. Topics include data confidentiality, data linkage, and national practices in policy areas such as public health, education and employment. It offers scholars as well as practitioners from public administrations, consultancy firms and nongovernmental organizations insights into counterfactual impact evaluation methods and the potential of data-based policy and program evaluation.

Linking Sensitive Data - Peter Christen 2021-10-19

This book provides modern technical answers to the legal requirements of pseudonymisation as recommended by privacy legislation. It covers topics such as modern regulatory frameworks for sharing and linking sensitive information, concepts and algorithms for privacy-preserving record linkage and their computational aspects, practical considerations such as dealing with dirty and missing data, as well as privacy, risk, and performance assessment measures. Existing techniques for

privacy-preserving record linkage are evaluated empirically and real-world application examples that scale to population sizes are described. The book also includes pointers to freely available software tools, benchmark data sets, and tools to generate synthetic data that can be used to test and evaluate linkage techniques. This book consists of fourteen chapters grouped into four parts, and two appendices. The first part introduces the reader to the topic of linking sensitive data, the second part covers methods and techniques to link such data, the third part discusses aspects of practical importance, and the fourth part provides an outlook of future challenges and open research problems relevant to linking sensitive databases. The appendices provide pointers and describe freely available, open-source software systems that allow the linkage of sensitive data, and provide further details about the evaluations presented. A companion Web site at https://dmm.anu.edu.au/lsdbook2020 provides additional material and Python programs used in the book. This book is mainly written for applied scientists, researchers, and advanced practitioners in governments, industry, and universities who are concerned with developing, implementing, and deploying systems and tools to share sensitive information in administrative, commercial, or medical databases. The Book describes how linkage methods work and how to evaluate their performance. It covers all the major concepts and methods and also discusses practical matters such as computational efficiency, which are critical if the methods are to be used in practice – and it does all this in a highly

accessible way!David J. Hand, Imperial College, London

Data Mining for Business Intelligence - Galit Shmueli 2006-12-11

Learn how to develop models for classification, prediction, and customer segmentation with the help of Data Mining for Business Intelligence In today's world, businesses are becoming more capable of accessing their ideal consumers, and an understanding of data mining contributes to this success. Data Mining for Business Intelligence, which was developed from a course taught at the Massachusetts Institute of Technology's Sloan School of Management, and the University of Maryland's Smith School of Business, uses real data and actual cases to illustrate the applicability of data mining intelligence to the development of successful business models. Featuring XLMiner, the Microsoft Office Excel add-in, this book allows readers to follow along and implement algorithms at their own speed, with a minimal learning curve. In addition, students and practitioners of data mining techniques are presented with hands-on, business-oriented applications. An abundant amount of exercises and examples are provided to motivate learning and understanding. Data Mining for Business Intelligence: Provides both a theoretical and practical understanding of the key methods of classification, prediction, reduction, exploration, and affinity analysis Features a business decision-making context for these key methods Illustrates the application and interpretation of these methods using real business cases and data This book helps readers understand the beneficial relationship that can be established between data mining and smart business practices, and is an

excellent learning tool for creating valuable strategies and making wiser business decisions.

Health Data in the Information Age - Institute of Medicine 1994-01-01

Regional health care databases are being established around the country with the goal of providing timely and useful information to policymakers, physicians, and patients. But their emergence is raising important and sometimes controversial questions about the collection, quality, and appropriate use of health care data. Based on experience with databases now in operation and in development, Health Data in the Information Age provides a clear set of guidelines and principles for exploiting the potential benefits of aggregated health dataâ€"without jeopardizing confidentiality. A panel of experts identifies characteristics of emerging health database organizations (HDOs). The committee explores how HDOs can maintain the quality of their data, what policies and practices they should adopt, how they can prepare for linkages with computer-based patient records, and how diverse groups from researchers to health care administrators might use aggregated data. Health Data in the Information Age offers frank analysis and guidelines that will be invaluable to anyone interested in the operation of health care databases.

Data Analytics in Medicine: Concepts, Methodologies, Tools, and Applications - Management Association, Information Resources 2019-12-06

Advancements in data science have created opportunities to sort, manage, and analyze large amounts of data more effectively and efficiently. Applying these new technologies to the healthcare industry, which has vast quantities of patient

and medical data and is increasingly becoming more data-reliant, is crucial for refining medical practices and patient care. Data Analytics in Medicine: Concepts, Methodologies, Tools, and Applications is a vital reference source that examines practical applications of healthcare analytics for improved patient care, resource allocation, and medical performance, as well as for diagnosing, predicting, and identifying at-risk populations. Highlighting a range of topics such as data security and privacy, health informatics, and predictive analytics, this multi-volume book is ideally designed for doctors, hospital administrators, nurses, medical professionals, IT specialists, computer engineers, information technologists, biomedical engineers, data-processing specialists, healthcare practitioners, academicians, and researchers interested in current research on the connections between data analytics in the field of medicine.

Quality Measures in Data Mining - Fabrice Guillet 2007-01-08

This book presents recent advances in quality measures in data mining.

Database Technologies: Concepts, Methodologies, Tools, and Applications - Erickson, John 2009-02-28

"This reference expands the field of database technologies through four-volumes of in-depth, advanced research articles from nearly 300 of the world's leading professionals"--Provided by publisher.

R for Data Science - Hadley Wickham 2016-12-12

Learn how to use R to turn raw data into insight, knowledge, and understanding. This book introduces you to R, RStudio, and the tidyverse, a collection of R packages designed to work together to make data science fast, fluent, and fun. Suitable for readers

with no previous programming experience, R for Data Science is designed to get you doing data science as quickly as possible. Authors Hadley Wickham and Garrett Grolemund guide you through the steps of importing, wrangling, exploring, and modeling your data and communicating the results. You'll get a complete, big-picture understanding of the data science cycle, along with basic tools you need to manage the details. Each section of the book is paired with exercises to help you practice what you've learned along the way. You'll learn how to: Wrangle—transform your datasets into a form convenient for analysis Program—learn powerful R tools for solving data problems with greater clarity and ease Explore—examine your data, generate hypotheses, and quickly test them Model—provide a low-dimensional summary that captures true "signals" in your dataset Communicate—learn R Markdown for integrating prose, code, and results

Data Matching - Peter Christen 2012-07-04 Data matching (also known as record or data linkage, entity resolution, object identification, or field matching) is the task of identifying, matching and merging records that correspond to the same entities from several databases or even within one database. Based on research in various domains including applied statistics, health informatics, data mining, machine learning, artificial intelligence, database management, and digital libraries, significant advances have been achieved over the last decade in all aspects of the data matching process, especially on how to improve the accuracy of data matching, and its scalability to large databases. Peter Christen’s book is divided into three parts: Part I, “Overview”, introduces

the subject by presenting several sample applications and their special challenges, as well as a general overview of a generic data matching process. Part II, “Steps of the Data Matching Process”, then details its main steps like pre-processing, indexing, field and record comparison, classification, and quality evaluation. Lastly, part III, “Further Topics”, deals with specific aspects like privacy, real-time matching, or matching unstructured data. Finally, it briefly describes the main features of many research and open source systems available today. By providing the reader with a broad range of data matching concepts and techniques and touching on all aspects of the data matching process, this book helps researchers as well as students specializing in data quality or data matching aspects to familiarize themselves with recent research advances and to identify open research challenges in the area of data matching. To this end, each chapter of the book includes a final section that provides pointers to further background and research material. Practitioners will better understand the current state of the art in data matching as well as the internal workings and limitations of current systems. Especially, they will learn that it is often not feasible to simply implement an existing off-the-shelf data matching system without substantial adaption and customization. Such practical considerations are discussed for each of the major steps in the data matching process.

Data Matching - Peter Christen 2014-08-09

Data matching (also known as record or data linkage, entity resolution, object identification, or field matching) is the task of identifying, matching and merging records that correspond to the same

entities from several databases or even within one database. Based on research in various domains including applied statistics, health informatics, data mining, machine learning, artificial intelligence, database management, and digital libraries, significant advances have been achieved over the last decade in all aspects of the data matching process, especially on how to improve the accuracy of data matching, and its scalability to large databases. Peter Christen's book is divided into three parts: Part I, "Overview", introduces the subject by presenting several sample applications and their special challenges, as well as a general overview of a generic data matching process. Part II, "Steps of the Data Matching Process", then details its main steps like pre-processing, indexing, field and record comparison, classification, and quality evaluation. Lastly, part III, "Further Topics", deals with specific aspects like privacy, real-time matching, or matching unstructured data. Finally, it briefly describes the main features of many research and open source systems available today. By providing the reader with a broad range of data matching concepts and techniques and touching on all aspects of the data matching process, this book helps researchers as well as students specializing in data quality or data matching aspects to familiarize themselves with recent research advances and to identify open research challenges in the area of data matching. To this end, each chapter of the book includes a final section that provides pointers to further background and research material. Practitioners will better understand the current state of the art in data matching as well as the internal workings

and limitations of current systems. Especially, they will learn that it is often not feasible to simply implement an existing off-the-shelf data matching system without substantial adaption and customization. Such practical considerations are discussed for each of the major steps in the data matching process.

Flexible Imputation of Missing Data, Second Edition - Stef van Buuren 2018-07-17 Missing data pose challenges to real-life data analysis. Simple ad-hoc fixes, like deletion or mean imputation, only work under highly restrictive conditions, which are often not met in practice. Multiple imputation replaces each missing value by multiple plausible values. The variability between these replacements reflects our ignorance of the true (but missing) value. Each of the completed data set is then analyzed by standard methods, and the results are pooled to obtain unbiased estimates with correct confidence intervals. Multiple imputation is a general approach that also inspires novel solutions to old problems by reformulating the task at hand as a missing-data problem. This is the second edition of a popular book on multiple imputation, focused on explaining the application of methods through detailed worked examples using the MICE package as developed by the author. This new edition incorporates the recent developments in this fast-moving field. This class-tested book avoids mathematical and technical details as much as possible: formulas are accompanied by verbal statements that explain the formula in accessible terms. The book sharpens the reader's intuition on how to think about missing data, and provides all the tools needed to execute a well-grounded

quantitative analysis in the presence of missing data.

Data Wrangling with Python - Jacqueline Kazil 2016-02-04

How do you take your data analysis skills beyond Excel to the next level? By learning just enough Python to get stuff done. This hands-on guide shows non-programmers like you how to process information that's initially too messy or difficult to access. You don't need to know a thing about the Python programming language to get started. Through various step-by-step exercises, you'll learn how to acquire, clean, analyze, and present data efficiently. You'll also discover how to automate your data process, schedule file- editing and clean-up tasks, process larger datasets, and create compelling stories with data you obtain. Quickly learn basic Python syntax, data types, and language concepts Work with both machine-readable and human-consumable data Scrape websites and APIs to find a bounty of useful information Clean and format data to eliminate duplicates and errors in your datasets Learn when to standardize data and when to test and script data cleanup Explore and analyze your datasets with new Python libraries and techniques Use Python solutions to automate your entire data-wrangling process

Real-time Linked Dataspaces - Edward Curry 2020-10-09

This open access book explores the dataspace paradigm as a best-effort approach to data management within data ecosystems. It establishes the theoretical foundations and principles of real-time linked dataspaces as a data platform for intelligent systems. The book introduces a set of specialized best-effort techniques and models to enable loose administrative proximity and semantic integration for managing and processing events and streams. The book is

divided into five major parts: Part I "Fundamentals and Concepts" details the motivation behind and core concepts of real-time linked dataspaces, and establishes the need to evolve data management techniques in order to meet the challenges of enabling data ecosystems for intelligent systems within smart environments. Further, it explains the fundamental concepts of dataspaces and the need for specialization in the processing of dynamic real-time data. Part II "Data Support Services" explores the design and evaluation of critical services, including catalog, entity management, query and search, data service discovery, and human-in-the-loop. In turn, Part III "Stream and Event Processing Services" addresses the design and evaluation of the specialized techniques created for real-time support services including complex event processing, event service composition, stream dissemination, stream matching, and approximate semantic matching. Part IV "Intelligent Systems and Applications" explores the use of real-time linked dataspaces within real-world smart environments. In closing, Part V "Future Directions" outlines future research challenges for dataspaces, data ecosystems, and intelligent systems. Readers will gain a detailed understanding of how the dataspace paradigm is now being used to enable data ecosystems for intelligent systems within smart environments. The book covers the fundamental theory, the creation of new techniques needed for support services, and lessons learned from real-world intelligent systems and applications focused on sustainability. Accordingly, it will benefit not only researchers and graduate students in the fields

of data management, big data, and IoT, but also professionals who need to create advanced data management platforms for intelligent systems, smart environments, and data ecosystems. This work was published by Saint Philip Street Press pursuant to a Creative Commons license permitting commercial use. All rights not granted by the work's license are retained by the author or authors.

Secondary Analysis of Electronic Health Records - MIT Critical Data 2016-09-09

This book trains the next generation of scientists representing different disciplines to leverage the data generated during routine patient care. It formulates a more complete lexicon of evidence-based recommendations and support shared, ethical decision making by doctors with their patients. Diagnostic and therapeutic technologies continue to evolve rapidly, and both individual practitioners and clinical teams face increasingly complex ethical decisions. Unfortunately, the current state of medical knowledge does not provide the guidance to make the majority of clinical decisions on the basis of evidence. The present research infrastructure is inefficient and frequently produces unreliable results that cannot be replicated. Even randomized controlled trials (RCTs), the traditional gold standards of the research reliability hierarchy, are not without limitations. They can be costly, labor intensive, and slow, and can return results that are seldom generalizable to every patient population. Furthermore, many pertinent but unresolved clinical and medical systems issues do not seem to have attracted the interest of the research enterprise, which has come to focus instead on cellular and molecular investigations and single-agent (e.g.,

a drug or device) effects. For clinicians, the end result is a bit of a “data desert” when it comes to making decisions. The new research infrastructure proposed in this book will help the medical profession to make ethically sound and well informed decisions for their patients.

Sharing Clinical Trial Data - Institute of Medicine 2015-04-20

Data sharing can accelerate new discoveries by avoiding duplicative trials, stimulating new ideas for research, and enabling the maximal scientific knowledge and benefits to be gained from the efforts of clinical trial participants and investigators. At the same time, sharing clinical trial data presents risks, burdens, and challenges. These include the need to protect the privacy and honor the consent of clinical trial participants; safeguard the legitimate economic interests of sponsors; and guard against invalid secondary analyses, which could undermine trust in clinical trials or otherwise harm public health. Sharing Clinical Trial Data presents activities and strategies for the responsible sharing of clinical trial data. With the goal of increasing scientific knowledge to lead to better therapies for patients, this book identifies guiding principles and makes recommendations to maximize the benefits and minimize risks. This report offers guidance on the types of clinical trial data available at different points in the process, the points in the process at which each type of data should be shared, methods for sharing data, what groups should have access to data, and future knowledge and infrastructure needs. Responsible sharing of clinical trial data will allow other investigators to replicate published

findings and carry out additional analyses, strengthen the evidence base for regulatory and clinical decisions, and increase the scientific knowledge gained from investments by the funders of clinical trials. The recommendations of Sharing Clinical Trial Data will be useful both now and well into the future as improved sharing of data leads to a stronger evidence base for treatment. This book will be of interest to stakeholders across the spectrum of research-- from funders, to researchers, to journals, to physicians, and ultimately, to patients.

The Algorithmic Foundations of Differential Privacy - Cynthia Dwork 2014

The problem of privacy-preserving data analysis has a long history spanning multiple disciplines. As electronic data about individuals becomes increasingly detailed, and as technology enables ever more powerful collection and curation of these data, the need increases for a robust, meaningful, and mathematically rigorous definition of privacy, together with a computationally rich class of algorithms that satisfy this definition. Differential Privacy is such a definition. The Algorithmic Foundations of Differential Privacy starts out by motivating and discussing the meaning of differential privacy, and proceeds to explore the fundamental techniques for achieving differential privacy, and the application of these techniques in creative combinations, using the query-release problem as an ongoing example. A key point is that, by rethinking the computational goal, one can often obtain far better results than would be achieved by methodically replacing each step of a non-private computation with a differentially private implementation. Despite some powerful computational results, there are still

fundamental limitations. Virtually all the algorithms discussed herein maintain differential privacy against adversaries of arbitrary computational power -- certain algorithms are computationally intensive, others are efficient. Computational complexity for the adversary and the algorithm are both discussed. The monograph then turns from fundamentals to applications other than query-release, discussing differentially private methods for mechanism design and machine learning. The vast majority of the literature on differentially private algorithms considers a single, static, database that is subject to many analyses. Differential privacy in other models, including distributed databases and computations on data streams, is discussed. The Algorithmic Foundations of Differential Privacy is meant as a thorough introduction to the problems and techniques of differential privacy, and is an invaluable reference for anyone with an interest in the topic.

Data Mining with Rattle and R - Graham Williams 2011-08-04

Data mining is the art and science of intelligent data analysis. By building knowledge from information, data mining adds considerable value to the ever increasing stores of electronic data that abound today. In performing data mining many decisions need to be made regarding the choice of methodology, the choice of data, the choice of tools, and the choice of algorithms. Throughout this book the reader is introduced to the basic concepts and some of the more popular algorithms of data mining. With a focus on the hands-on end-to-end process for data mining, Williams guides the reader through various capabilities of the easy

to use, free, and open source Rattle Data Mining Software built on the sophisticated R Statistical Software. The focus on doing data mining rather than just reading about data mining is refreshing. The book covers data understanding, data preparation, data refinement, model building, model evaluation, and practical deployment. The reader will learn to rapidly deliver a data mining project using software easily installed for free from the Internet. Coupling Rattle with R delivers a very sophisticated data mining environment with all the power, and more, of the many commercial offerings.

Fixing Access Annoyances - Phil Mitchell 2006-02-21

Provides a collection of tips on fixing annoyances found in Microsoft Access, covering such topics as performance, security, database design, queries, forms, page layout, macros, and expressions.

Entity Resolution and Information Quality - John R. Talburt 2011-01-14

Entity Resolution and Information Quality presents topics and definitions, and clarifies confusing terminologies regarding entity resolution and information quality. It takes a very wide view of IQ, including its six-domain framework and the skills formed by the International Association for Information and Data Quality {IAIDQ). The book includes chapters that cover the principles of entity resolution and the principles of Information Quality, in addition to their concepts and terminology. It also discusses the Fellegi-Sunter theory of record linkage, the Stanford Entity Resolution Framework, and the Algebraic Model for Entity Resolution, which are the major theoretical models that support Entity

Resolution. In relation to this, the book briefly discusses entity-based data integration (EBDI) and its model, which serve as an extension of the Algebraic Model for Entity Resolution. There is also an explanation of how the three commercial ER systems operate and a description of the non-commercial open-source system known as OYSTER. The book concludes by discussing trends in entity resolution research and practice. Students taking IT courses and IT professionals will find this book invaluable. First authoritative reference explaining entity resolution and how to use it effectively Provides practical system design advice to help you get a competitive advantage Includes a companion site with synthetic customer data for applicatory exercises, and access to a Java-based Entity Resolution program.

Data Quality and Record Linkage Techniques - Thomas N. Herzog 2007-05-23

This book offers a practical understanding of issues involved in improving data quality through editing, imputation, and record linkage. The first part of the book deals with methods and models, focusing on the Fellegi-Holt edit-imputation model, the Little-Rubin multiple-imputation scheme, and the Fellegi-Sunter record linkage model. The second part presents case studies in which these techniques are applied in a variety of areas, including mortgage guarantee insurance, medical, biomedical, highway safety, and social insurance as well as the construction of list frames and administrative lists. This book offers a mixture of practical advice, mathematical rigor, management insight and philosophy.

Computers at Risk - National Research Council 1990-02-01

Computers at Risk

presents a comprehensive agenda for developing nationwide policies and practices for computer security. Specific recommendations are provided for industry and for government agencies engaged in computer security activities. The volume also outlines problems and opportunities in computer security research, recommends ways to improve the research infrastructure, and suggests topics for investigators. The book explores the diversity of the field, the need to engineer countermeasures based on speculation of what experts think computer attackers may do next, why the technology community has failed to respond to the need for enhanced security systems, how innovators could be encouraged to bring more options to the marketplace, and balancing the importance of security against the right of privacy.

Data Deduplication Approaches - Tin Thein Thwel 2020-11-25

In the age of data science, the rapidly increasing amount of data is a major concern in numerous applications of computing operations and data storage. Duplicated data or redundant data is a main challenge in the field of data science research. Data Deduplication Approaches: Concepts, Strategies, and Challenges shows readers the various methods that can be used to eliminate multiple copies of the same files as well as duplicated segments or chunks of data within the associated files. Due to ever-increasing data duplication, its deduplication has become an especially useful field of research for storage environments, in particular persistent data storage. Data Deduplication Approaches provides readers with an overview of the concepts and background of data deduplication approaches, then proceeds to demonstrate in technical detail the strategies and

challenges of real-time implementations of handling big data, data science, data backup, and recovery. The book also includes future research directions, case studies, and real-world applications of data deduplication, focusing on reduced storage, backup, recovery, and reliability. Includes data deduplication methods for a wide variety of applications Includes concepts and implementation strategies that will help the reader to use the suggested methods Provides a robust set of methods that will help readers to appropriately and judiciously use the suitable methods for their applications Focuses on reduced storage, backup, recovery, and reliability, which are the most important aspects of implementing data deduplication approaches Includes case studies

Data Quality - Carlo Batini 2006-09-27

Poor data quality can seriously hinder or damage the efficiency and effectiveness of organizations and businesses. The growing awareness of such repercussions has led to major public initiatives like the "Data Quality Act" in the USA and the "European 2003/98" directive of the European Parliament. Batini and Scannapieco present a comprehensive and systematic introduction to the wide set of issues related to data quality. They start with a detailed description of different data quality dimensions, like accuracy, completeness, and consistency, and their importance in different types of data, like federated data, web data, or time-dependent data, and in different data categories classified according to frequency of change, like stable, long-term, and frequently changing data. The book's extensive description of techniques and methodologies from core data quality research as

well as from related fields like data mining, probability theory, statistical data analysis, and machine learning gives an excellent overview of the current state of the art. The presentation is completed by a short description and critical comparison of tools and practical methodologies, which will help readers to resolve their own quality problems. This book is an ideal combination of the soundness of theoretical foundations and the applicability of practical approaches. It is ideally suited for everyone – researchers, students, or professionals – interested in a comprehensive overview of data quality issues. In addition, it will serve as the basis for an introductory course or for self-study on this topic.

Mining of Massive Datasets – Jure Leskovec 2014-11-13

Now in its second edition, this book focuses on practical algorithms for mining data from even the largest datasets.

Advanced Data Mining Techniques – David L. Olson 2008-01-01

This book covers the fundamental concepts of data mining, to demonstrate the potential of gathering large sets of data, and analyzing these data sets to gain useful business understanding. The book is organized in three parts. Part I introduces concepts. Part II describes and demonstrates basic data mining algorithms. It also contains chapters on a number of different techniques often used in data mining. Part III focuses on business applications of data mining.

Research Anthology on Privatizing and Securing Data – Management Association, Information Resources 2021-04-23

With the immense amount of data that is now available online, security concerns have been an issue from the start, and have grown as new technologies are

increasingly integrated in data collection, storage, and transmission. Online cyber threats, cyber terrorism, hacking, and other cybercrimes have begun to take advantage of this information that can be easily accessed if not properly handled. New privacy and security measures have been developed to address this cause for concern and have become an essential area of research within the past few years and into the foreseeable future. The ways in which data is secured and privatized should be discussed in terms of the technologies being used, the methods and models for security that have been developed, and the ways in which risks can be detected, analyzed, and mitigated. The Research Anthology on Privatizing and Securing Data reveals the latest tools and technologies for privatizing and securing data across different technologies and industries. It takes a deeper dive into both risk detection and mitigation, including an analysis of cybercrimes and cyber threats, along with a sharper focus on the technologies and methods being actively implemented and utilized to secure data online. Highlighted topics include information governance and privacy, cybersecurity, data protection, challenges in big data, security threats, and more. This book is essential for data analysts, cybersecurity professionals, data scientists, security analysts, IT specialists, practitioners, researchers, academicians, and students interested in the latest trends and technologies for privatizing and securing data.

The Behavioral and Social Sciences - National Research Council 1988-02-01

This volume explores the scientific frontiers and leading edges of research across the fields of anthropology,

economics, political science, psychology, sociology, history, business, education, geography, law, and psychiatry, as well as the newer, more specialized areas of artificial intelligence, child development, cognitive science, communications, demography, linguistics, and management and decision science. It includes recommendations concerning new resources, facilities, and programs that may be needed over the next several years to ensure rapid progress and provide a high level of returns to basic research.

SQL Cookbook - Anthony Molinaro 2006

A guide to SQL covers such topics as retrieving records, metadata queries, working with strings, data arithmetic, date manipulation, reporting and warehousing, and hierarchical queries.

Registries for Evaluating Patient Outcomes - Agency for Healthcare Research and Quality/AHRQ 2014-04-01

This User's Guide is intended to support the design, implementation, analysis, interpretation, and quality evaluation of registries created to increase understanding of patient outcomes. For the purposes of this guide, a patient registry is an organized system that uses observational study methods to collect uniform data (clinical and other) to evaluate specified outcomes for a population defined by a particular disease, condition, or exposure, and that serves one or more predetermined scientific, clinical, or policy purposes. A registry database is a file (or files) derived from the registry. Although registries can serve many purposes, this guide focuses on registries created for one or more of the following purposes: to describe the natural history of disease, to determine clinical effectiveness or cost-effectiveness of health

care products and services, to measure or monitor safety and harm, and/or to measure quality of care. Registries are classified according to how their populations are defined. For example, product registries include patients who have been exposed to biopharmaceutical products or medical devices. Health services registries consist of patients who have had a common procedure, clinical encounter, or hospitalization. Disease or condition registries are defined by patients having the same diagnosis, such as cystic fibrosis or heart failure. The User's Guide was created by researchers affiliated with AHRQ's Effective Health Care Program, particularly those who participated in AHRQ's DEcIDE (Developing Evidence to Inform Decisions About Effectiveness) program. Chapters were subject to multiple internal and external independent reviews.

The Semantic Web. Latest Advances and New Domains – Harald Sack 2016-05-22

The 47 revised full papers presented together with three invited talks were carefully reviewed and selected from 204 submissions. This program was completed by a demonstration and poster session, in which researchers had the chance to present their latest results and advances in the form of live demos. In addition, the PhD Symposium program included 10 contributions, selected out of 21 submissions. The core tracks of the research conference were complemented with new tracks focusing on linked data; machine learning; mobile web, sensors and semantic streams; natural language processing and information retrieval; reasoning; semantic data management, big data, and scalability; services, APIs, processes and cloud computing; smart cities, urban and geospatial

data; trust and privacy; and vocabularies, schemas, and ontologies.

Methodological Developments in Data Linkage - Katie Harron 2015-09-22

A comprehensive compilation of new developments in data linkage methodology The increasing availability of large administrative databases has led to a dramatic rise in the use of data linkage, yet the standard texts on linkage are still those which describe the seminal work from the 1950-60s, with some updates. Linkage and analysis of data across sources remains problematic due to lack of discriminatory and accurate identifiers, missing data and regulatory issues. Recent developments in data linkage methodology have concentrated on bias and analysis of linked data, novel approaches to organising relationships between databases and privacy-preserving linkage. Methodological Developments in Data Linkage brings together a collection of contributions from members of the international data linkage community, covering cutting edge methodology in this field. It presents opportunities and challenges provided by linkage of large and often complex datasets, including analysis problems, legal and security aspects, models for data access and the development of novel research areas. New methods for handling uncertainty in analysis of linked data, solutions for anonymised linkage and alternative models for data collection are also discussed. Key Features: Presents cutting edge methods for a topic of increasing importance to a wide range of research areas, with applications to data linkage systems internationally Covers the essential issues associated with data linkage today Includes examples based on real data linkage systems, highlighting the

opportunities, successes and challenges that the increasing availability of linkage data provides Novel approach incorporates technical aspects of both linkage, management and analysis of linked data This book will be of core interest to academics, government employees, data holders, data managers, analysts and statisticians who use administrative data. It will also appeal to researchers in a variety of areas, including epidemiology, biostatistics, social statistics, informatics, policy and public health.

Semantic Systems. In the Era of Knowledge Graphs - Eva Blomqvist 2020-10-26

This open access book constitutes the refereed proceedings of the 16th International Conference on Semantic Systems, SEMANTiCS 2020, held in Amsterdam, The Netherlands, in September 2020. The conference was held virtually due to the COVID-19 pandemic.

Data Mining: Concepts and Techniques - Jiawei Han 2011-06-09

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

The Book of R - Tilman M. Davies 2016-07-16

The Book of R is a comprehensive, beginner-friendly guide to R, the world's most popular programming language for statistical analysis. Even if you have no programming experience and little more than a grounding in the basics of mathematics, you'll find everything you need to begin using R effectively for statistical analysis. You'll start with the basics, like how to handle data and write simple programs, before moving on to more advanced topics, like producing statistical summaries of your data and performing statistical tests and modeling. You'll even learn how to create impressive data visualizations with R's basic graphics tools and contributed packages, like ggplot2 and ggvis, as well as interactive 3D visualizations using the rgl package. Dozens

of hands-on exercises (with downloadable solutions) take you from theory to practice, as you learn: –The fundamentals of programming in R, including how to write data frames, create functions, and use variables, statements, and loops –Statistical concepts like exploratory data analysis, probabilities, hypothesis tests, and regression modeling, and how to execute them in R –How to access R's thousands of functions, libraries, and data sets –How to draw valid and useful conclusions from your data –How to create publication-quality graphics of your results Combining detailed explanations with real-world examples and exercises, this book will provide you with a solid understanding of both statistics and the depth of R's functionality. Make The Book of R your doorway into the growing world of data analysis.

Strengthening Forensic Science in the United States - National Research Council 2009-07-29

Scores of talented and dedicated people serve the forensic science community, performing vitally important work. However, they are often constrained by lack of adequate resources, sound policies, and national support. It is clear that change and advancements, both systematic and scientific, are needed in a number of forensic science disciplines to ensure the reliability of work, establish enforceable standards, and promote best practices with consistent application. Strengthening Forensic Science in the United States: A Path Forward provides a detailed plan for addressing these needs and suggests the creation of a new government entity, the National Institute of Forensic Science, to establish and enforce standards within the forensic science community. The benefits of improving and

regulating the forensic science disciplines are clear: assisting law enforcement officials, enhancing homeland security, and reducing the risk of wrongful conviction and exoneration. Strengthening Forensic Science in the United States gives a full account of what is needed to advance the forensic science disciplines, including upgrading of systems and organizational structures, better training, widespread adoption of uniform and enforceable best practices, and mandatory certification and accreditation programs. While this book provides an essential call-to-action for congress and policy makers, it also serves as a vital tool for law enforcement agencies, criminal prosecutors and attorneys, and forensic science educators.