

14 Predicate Logic Stanford Infolab

As recognized, adventure as with ease as experience just about lesson, amusement, as well as conformity can be gotten by just checking out a books **14 predicate logic stanford infolab** afterward it is not directly done, you could agree to even more vis--vis this life, re the world.

We present you this proper as with ease as simple showing off to acquire those all. We allow 14 predicate logic stanford infolab and numerous books collections from fictions to scientific research in any way. among them is this 14 predicate logic stanford infolab that can be your partner.

Linked Open Data -- Creating Knowledge Out of Interlinked Data - Sören Auer

2014-07-31

Linked Open Data (LOD) is a pragmatic approach for realizing the Semantic Web vision of making the Web a global, distributed, semantics-based information system. This book presents an overview on the results of the research project "LOD2 -- Creating Knowledge

out of Interlinked Data". LOD2 is a large-scale integrating project co-funded by the European Commission within the FP7 Information and Communication Technologies Work Program. Commencing in September 2010, this 4-year project comprised leading Linked Open Data research groups, companies, and service providers from across 11 European countries and South Korea. The aim of this project was to

advance the state-of-the-art in research and development in four key areas relevant for Linked Data, namely 1. RDF data management; 2. the extraction, creation, and enrichment of structured RDF data; 3. the interlinking and fusion of Linked Data from different sources and 4. the authoring, exploration and visualization of Linked Data.

Towards Semantically Enabled Complex Event Processing Robin Keskisärkkä 2017-10-30

The Semantic Web provides a framework for semantically annotating data on the web, and the Resource Description Framework (RDF) supports the integration of structured data represented in heterogeneous formats. Traditionally, the Semantic Web has focused primarily on more or less static data, but information on the web today is becoming increasingly dynamic. RDF Stream Processing (RSP) systems address this issue by adding support for streaming data and continuous query processing. To some extent, RSP systems can be

used to perform complex event processing (CEP), where meaningful high-level events are generated based on low-level events from multiple sources; however, there are several challenges with respect to using RSP in this context. Event models designed to represent static event information lack several features required for CEP, and are typically not well suited for stream reasoning. The dynamic nature of streaming data also greatly complicates the development and validation of RSP queries. Therefore, reusing queries that have been prepared ahead of time is important to be able to support real-time decision-making. Additionally, there are limitations in existing RSP implementations in terms of both scalability and expressiveness, where some features required in CEP are not supported by any of the current systems. The goal of this thesis work has been to address some of these challenges and the main contributions of the thesis are: (1) an event model ontology targeted at supporting

CEP; (2) a model for representing parameterized RSP queries as reusable templates; and (3) an architecture that allows RSP systems to be integrated for use in CEP. The proposed event model tackles issues specifically related to event modeling in CEP that have not been sufficiently covered by other event models, includes support for event encapsulation and event payloads, and can easily be extended to fit specific use-cases. The model for representing RSP query templates was designed as an extension to SPIN, a vocabulary that supports modeling of SPARQL queries as RDF. The extended model supports the current version of the RSP Query Language (RSP-QL) developed by the RDF Stream Processing Community Group, along with some of the most popular RSP query languages. Finally, the proposed architecture views RSP queries as individual event processing agents in a more general CEP framework. Additional event processing components can be integrated to provide support for operations that are not

supported in RSP, or to provide more efficient processing for specific tasks. We demonstrate the architecture in implementations for scenarios related to traffic-incident monitoring, criminal-activity monitoring, and electronic healthcare monitoring.

Senant i cs- Ori ent ed Nat ural Language Processi ng Vladimir Fomichov A. 2009-12-01
Gluecklich, die wissen, dass hinter allen Sprachen das Unsaegliche steht. Those are happy who know that behind all languages there is something unsaid Rainer Maria Rilke This book shows in a new way that a solution to a fundamental problem from one scienti?c ?eld can help to ?nd the solutions to important problems emerged in several other ?elds of science and technology. In modern science, the term "Natural Language" denotes the collection of all such languages that every language is used as a primary means of communication by people belonging to any country or any region. So Natural Language (NL) includes, in particular,

the English, Russian, and German languages. The applied computer systems processing natural language printed or written texts (NL-texts) or oral speech with respect to the fact that the words are associated with some meanings are called semantics-oriented natural language processing systems (NLPs). On one hand, this book is a snapshot of the current stage of a research program started many years ago and called Integral Formal Semantics (IFS) of NL. The goal of this program has been to develop the formal models and methods helping to overcome the difficulties of logical character associated with the engineering of semantics-oriented NLPs. The designers of such systems of arbitrary kinds will find in this book the formal means and algorithms being of great help in their work.

Large-Scale Graph Processing Using Apache Giraph - Sherif Sakr 2017-01-05

This book takes its reader on a journey through Apache Giraph, a popular distributed graph

processing platform designed to bring the power of big data processing to graph data. Designed as a step-by-step self-study guide for everyone interested in large-scale graph processing, it describes the fundamental abstractions of the system, its programming models and various techniques for using the system to process graph data at scale, including the implementation of several popular and advanced graph analytics algorithms. The book is organized as follows: Chapter 1 starts by providing a general background of the big data phenomenon and a general introduction to the Apache Giraph system, its abstraction, programming model and design architecture. Next, chapter 2 focuses on Giraph as a platform and how to use it. Based on a sample job, even more advanced topics like monitoring the Giraph application lifecycle and different methods for monitoring Giraph jobs are explained. Chapter 3 then provides an introduction to Giraph programming, introduces the basic Giraph graph model and explains how

to write Giraph programs. In turn, Chapter 4 discusses in detail the implementation of some popular graph algorithms including PageRank, connected components, shortest paths and triangle closing. Chapter 5 focuses on advanced Giraph programming, discussing common Giraph algorithmic optimizations, tunable Giraph configurations that determine the system's utilization of the underlying resources, and how to write a custom graph input and output format. Lastly, chapter 6 highlights two systems that have been introduced to tackle the challenge of large scale graph processing, GraphX and GraphLab, and explains the main commonalities and differences between these systems and Apache Giraph. This book serves as an essential reference guide for students, researchers and practitioners in the domain of large scale graph processing. It offers step-by-step guidance, with several code examples and the complete source code available in the related github repository. Students will find a comprehensive introduction

to and hands-on practice with tackling large scale graph processing problems using the Apache Giraph system, while researchers will discover thorough coverage of the emerging and ongoing advancements in big graph processing systems.

A Semantic Web Primer, third edition -
Grigoris Antoniou 2012-09-07

A new edition of the widely used guide to the key ideas, languages, and technologies of the Semantic Web The development of the Semantic Web, with machine-readable content, has the potential to revolutionize the World Wide Web and its uses. A Semantic Web Primer provides an introduction and guide to this continuously evolving field, describing its key ideas, languages, and technologies. Suitable for use as a textbook or for independent study by professionals, it concentrates on undergraduate-level fundamental concepts and techniques that will enable readers to proceed with building applications on their own and includes exercises,

project descriptions, and annotated references to relevant online materials. The third edition of this widely used text has been thoroughly updated, with significant new material that reflects a rapidly developing field. Treatment of the different languages (OWL2, rules) expands the coverage of RDF and OWL, defining the data model independently of XML and including coverage of N3/Turtle and RDFa. A chapter is devoted to OWL2, the new W3C standard. This edition also features additional coverage of the query language SPARQL, the rule language RIF and the possibility of interaction between rules and ontology languages and applications. The chapter on Semantic Web applications reflects the rapid developments of the past few years. A new chapter offers ideas for term projects. Additional material, including updates on the technological trends and research directions, can be found at <http://www.semanticwebprimer.org>.

Security Protocols XXVII - Jonathan Anderson

2020-08-20

The volume LNCS 12287 constitutes the proceedings of the 27th International Workshop on Security Protocols, held in Cambridge, UK, in April 2019. The volume consists of 16 thoroughly revised invited papers presented together with the respective transcripts of discussions. The theme of this year's workshop was "Security Protocols for Humans" The topics covered included Designing for Humans and Understanding Humans, Human Limitations in Security, Secure sharing and collaboration and much more.

Foundations of Computer Science - Alfred V. Aho
1994-10-15

Natural Language Processing and Information Systems - Chris Biemann 2015-06-03

This book constitutes the refereed proceedings of the 20th International Conference on Applications of Natural Language to Information Systems, NLDB 2015, held in Passau, Germany,

Downloaded from constructivworks.com
on by guest

in June 2015. The 18 full papers, 15 short papers, 14 poster and demonstration papers presented were carefully reviewed and selected from 100 submissions. The papers cover the following topics: information extraction, distributional semantics, querying and question answering systems, context-aware NLP, cognitive and semantic computing, sentiment and opinion analysis, information extraction and social media, NLP and usability, text classification and extraction, and posters and demonstrations.

Foundations of Fuzzy Logic and Semantic Web Languages (Open Access) - Umberto Straccia
2016-04-19

Managing vagueness/fuzziness is starting to play an important role in Semantic Web research, with a large number of research efforts underway. Foundations of Fuzzy Logic and Semantic Web Languages provides a rigorous and succinct account of the mathematical methods and tools used for representing and

reasoning with fuzzy information within Semantic

Privacy-Preserving Data Publishing Bige-Chung Chen
2009-10-14

This book is dedicated to those who have something to hide. It is a book about "privacy preserving data publishing" -- the art of publishing sensitive personal data, collected from a group of individuals, in a form that does not violate their privacy. This problem has numerous and diverse areas of application, including releasing Census data, search logs, medical records, and interactions on a social network. The purpose of this book is to provide a detailed overview of the current state of the art as well as open challenges, focusing particular attention on four key themes: RIGOROUS PRIVACY POLICIES Repeated and highly-publicized attacks on published data have demonstrated that simplistic approaches to data publishing do not work. Significant recent advances have exposed the shortcomings of

naive (and not-so-naive) techniques. They have also led to the development of mathematically rigorous definitions of privacy that publishing techniques must satisfy; METRICS FOR DATA UTILITY While it is necessary to enforce stringent privacy policies, it is equally important to ensure that the published version of the data is useful for its intended purpose. The authors provide an overview of diverse approaches to measuring data utility; ENFORCEMENT MECHANISMS This book describes in detail various key data publishing mechanisms that guarantee privacy and utility; EMERGING APPLICATIONS The problem of privacy-preserving data publishing arises in diverse application domains with unique privacy and utility requirements. The authors elaborate on the merits and limitations of existing solutions, based on which we expect to see many advances in years to come.

Robust Stream Reasoning Under Uncertainty - Daniel de Leng 2019-11-08

Vast amounts of data are continually being generated by a wide variety of data producers. This data ranges from quantitative sensor observations produced by robot systems to complex unstructured human-generated texts on social media. With data being so abundant, the ability to make sense of these streams of data through reasoning is of great importance. Reasoning over streams is particularly relevant for autonomous robotic systems that operate in physical environments. They commonly observe this environment through incremental observations, gradually refining information about their surroundings. This makes robust management of streaming data and their refinement an important problem. Many contemporary approaches to stream reasoning focus on the issue of querying data streams in order to generate higher-level information by relying on well-known database approaches. Other approaches apply logic-based reasoning techniques, which rarely consider the

provenance of their symbolic interpretations. In this work, we integrate techniques for logic-based stream reasoning with the adaptive generation of the state streams needed to do the reasoning over. This combination deals with both the challenge of reasoning over uncertain streaming data and the problem of robustly managing streaming data and their refinement. The main contributions of this work are (1) a logic-based temporal reasoning technique based on path checking under uncertainty that combines temporal reasoning with qualitative spatial reasoning; (2) an adaptive reconfiguration procedure for generating and maintaining a data stream required to perform spatio-temporal stream reasoning over; and (3) integration of these two techniques into a stream reasoning framework. The proposed spatio-temporal stream reasoning technique is able to reason with intertemporal spatial relations by leveraging landmarks. Adaptive state stream generation allows the framework to

adapt to situations in which the set of available streaming resources changes. Management of streaming resources is formalised in the DyKnow model, which introduces a configuration life-cycle to adaptively generate state streams. The DyKnow-ROS stream reasoning framework is a concrete realisation of this model that extends the Robot Operating System (ROS). DyKnow-ROS has been deployed on the SoftBank Robotics NAO platform to demonstrate the system's capabilities in a case study on run-time adaptive reconfiguration. The results show that the proposed system - by combining reasoning over and reasoning about streams - can robustly perform stream reasoning, even when the availability of streaming resources changes.

The Quest for Artificial Intelligence - Nils J. Nilsson 2009-10-30

Artificial intelligence (AI) is a field within computer science that is attempting to build enhanced intelligence into computer systems. This book traces the history of the subject, from

the early dreams of eighteenth-century (and earlier) pioneers to the more successful work of today's AI engineers. AI is becoming more and more a part of everyone's life. The technology is already embedded in face-recognizing cameras, speech-recognition software, Internet search engines, and health-care robots, among other applications. The book's many diagrams and easy-to-understand descriptions of AI programs will help the casual reader gain an understanding of how these and other AI systems actually work. Its thorough (but unobtrusive) end-of-chapter notes containing citations to important source materials will be of great use to AI scholars and researchers. This book promises to be the definitive history of a field that has captivated the imaginations of scientists, philosophers, and writers for centuries.

Readings in Database Systems - Joseph M. Hellerstein 2005

The latest edition of a popular text and reference

on database research, with substantial new material and revision; covers classical literature and recent hot topics. Lessons from database research have been applied in academic fields ranging from bioinformatics to next-generation Internet architecture and in industrial uses including Web-based e-commerce and search engines. The core ideas in the field have become increasingly influential. This text provides both students and professionals with a grounding in database research and a technical context for understanding recent innovations in the field. The readings included treat the most important issues in the database area--the basic material for any DBMS professional. This fourth edition has been substantially updated and revised, with 21 of the 48 papers new to the edition, four of them published for the first time. Many of the sections have been newly organized, and each section includes a new or substantially revised introduction that discusses the context, motivation, and controversies in a particular

area, placing it in the broader perspective of database research. Two introductory articles, never before published, provide an organized, current introduction to basic knowledge of the field; one discusses the history of data models and query languages and the other offers an architectural overview of a database system. The remaining articles range from the classical literature on database research to treatments of current hot topics, including a paper on search engine architecture and a paper on application servers, both written expressly for this edition. The result is a collection of papers that are seminal and also accessible to a reader who has a basic familiarity with database systems.

Rule - Based Modeling and Computing on the Semantic Web - Monica Palmirani

2011-10-14

This book constitutes the refereed proceedings of the International RuleML Symposium, RuleML 2011-America, held in Fort Lauderdale, FL, USA, in November 2011 - collocated with the 22nd

International Joint Conference on Artificial Intelligence, IJCAI 2011. It is the second of two RuleML events that take place in 2011. The first RuleML Symposium, RuleML 2011-Europe, has been held in Barcelona, Spain, in July 2011. The 12 full papers, 5 short papers and 5 invited track and position papers presented together with 3 keynote speeches were carefully reviewed and selected from numerous submissions. The accepted papers address a wide range of rules, semantic technology, and cross-industry standards, rules and automated reasoning, rule-based event processing and reaction rules, vocabularies, ontologies and business rules, cloud computing and rules, clinical semantics and rules.

Computer Information Systems and Industrial Management - Khalid Saeed 2017-05-16

This book constitutes the proceedings of the 16th IFIP TC8 International Conference on Computer Information Systems and Industrial Management, CISIM 2017, held in Bialystok,

Poland, in June 2017. The 60 regular papers presented together with 5 keynotes were carefully reviewed and Selected from 85 submissions. They are organized in the following topical sections: algorithms; biometrics and pattern recognition applications; data analysis and information retrieval; engineering of enterprise software products; industrial management and other applications; modelling and optimization; various aspects of computer security.

Geospatial Data Science Techniques and Applications - Hassan A. Karimi 2017-10-24

Data science has recently gained much attention for a number of reasons, and among them is Big Data. Scientists (from almost all disciplines including physics, chemistry, biology, sociology, among others) and engineers (from all fields including civil, environmental, chemical, mechanical, among others) are faced with challenges posed by data volume, variety, and velocity, or Big Data. This book is designed to

highlight the unique characteristics of geospatial data, demonstrate the need to different approaches and techniques for obtaining new knowledge from raw geospatial data, and present select state-of-the-art geospatial data science techniques and how they are applied to various geoscience problems.

Big Data 2.0 Processing Systems - Sherif Sakr 2016-08-24

This book provides readers the “big picture” and a comprehensive survey of the domain of big data processing systems. For the past decade, the Hadoop framework has dominated the world of big data processing, yet recently academia and industry have started to recognize its limitations in several application domains and big data processing scenarios such as the large-scale processing of structured data, graph data and streaming data. Thus, it is now gradually being replaced by a collection of engines that are dedicated to specific verticals (e.g. structured data, graph data, and streaming

data). The book explores this new wave of systems, which it refers to as Big Data 2.0 processing systems. After Chapter 1 presents the general background of the big data phenomena, Chapter 2 provides an overview of various general-purpose big data processing systems that allow their users to develop various big data processing jobs for different application domains. In turn, Chapter 3 examines various systems that have been introduced to support the SQL flavor on top of the Hadoop infrastructure and provide competing and scalable performance in the processing of large-scale structured data. Chapter 4 discusses several systems that have been designed to tackle the problem of large-scale graph processing, while the main focus of Chapter 5 is on several systems that have been designed to provide scalable solutions for processing big data streams, and on other sets of systems that have been introduced to support the development of data pipelines between various

types of big data processing jobs and systems. Lastly, Chapter 6 shares conclusions and an outlook on future research challenges. Overall, the book offers a valuable reference guide for students, researchers and professionals in the domain of big data processing systems. Further, its comprehensive content will hopefully encourage readers to pursue further research on the subject.

On the Move to Meaningful Internet Systems: OTM 2013 Workshops - Yan Tang Demey 2013-09-25

This volume constitutes the refereed proceedings of the international workshops, Confederated International Workshops: OTM Academy, OTM Industry Case Studies Program, ACM, EI2N, ISDE, META4eS, ORM, SeDeS, SINCOM, SMS and SOMOCO 2013, held as part of OTM 2013 in Graz, Austria, in September 2013. The 75 revised full papers presented together with 12 posters and 5 keynote were carefully reviewed and selected from a total of

131 submissions. The papers are organized in topical sections on: On The Move Academy; Industry Case Studies Program; Adaptive Case Management and other non-workflow approaches to BPM; Enterprise Integration, Interoperability and Networking; Information Systems in Distributed Environment; Methods, Evaluation, Tools and Applications for the Creation and Consumption of Structured Data for the e-Society; Fact-Oriented Modeling; Semantics and Decision Making; Social Media Semantics; Social and Mobile Computing for collaborative environments; cooperative information systems; Ontologies, Data Bases and Applications of Semantics.

ALGORITHMS OF THE INTELLIGENT WEB -

Haralambos Marmanis 2011-03-01

Special Features: Learning Elements:· How to create recommendations just like those on Netflix and Amazon· How to implement Google's Pagerank algorithm· How to discover matches on social-networking sites· How to organize the

discussions on your favorite news group· How to select topics of interest from shared bookmarks· How to leverage user clicks· How to categorize emails based on their content· How to build applications that do targeted advertising· How to implement fraud detection About The Book: Algorithms of the Intelligent Web is an example-driven blueprint for creating applications that collect, analyze, and act on the massive quantities of data users leave in their wake as they use the web. You'll learn how to build Amazon- and Netflix-style recommendation engines, and how the same techniques apply to people matches on social-networking sites. See how click-trace analysis can result in smarter ad rotations. With a plethora of examples and extensive detail, this book shows you how to build Web 2.0 applications that are as smart as your users.

Engi nes of Order Bernhard Rieder 2020-05-18
Software has become a key component of contemporary life and algorithmic techniques

Downloaded from constructivworks.com
on by guest

that rank, classify, or recommend anything that fits into digital form are everywhere. This book approaches the field of information ordering conceptually as well as historically. Building on the philosophy of Gilbert Simondon and the cultural techniques tradition, it first examines the constructive and cumulative character of software and shows how software-making constantly draws on large reservoirs of existing knowledge and techniques. It then reconstructs the historical trajectories of a series of algorithmic techniques that have indeed become the building blocks for contemporary practices of ordering. Developed in opposition to centuries of library tradition, coordinate indexing, text processing, machine learning, and network algorithms instantiate dynamic, perspectivist, and interested forms of arranging information, ideas, or people. Embedded in technical infrastructures and economic logics, these techniques have become engines of order that transform the spaces they act upon.

Adding Sense - Mary Kalantzis 2020-05-07

Through a wide range of examples, from literature to social media, the book explores how meaning and communication interact.

Exploiting Semantic Web Knowledge Graphs in Data Mining - P. Ristoski 2019-06-28

Data Mining and Knowledge Discovery in Databases (KDD) is a research field concerned with deriving higher-level insights from data. The tasks performed in this field are knowledge intensive and can benefit from additional knowledge from various sources, so many approaches have been proposed that combine Semantic Web data with the data mining and knowledge discovery process. This book, Exploiting Semantic Web Knowledge Graphs in Data Mining, aims to show that Semantic Web knowledge graphs are useful for generating valuable data mining features that can be used in various data mining tasks. In Part I, Mining Semantic Web Knowledge Graphs, the author evaluates unsupervised feature generation

strategies from types and relations in knowledge graphs used in different data mining tasks such as classification, regression, and outlier detection. Part II, Semantic Web Knowledge Graphs Embeddings, proposes an approach that circumvents the shortcomings introduced with the approaches in Part I, developing an approach that is able to embed complete Semantic Web knowledge graphs in a low dimensional feature space where each entity and relation in the knowledge graph is represented as a numerical vector. Finally, Part III, Applications of Semantic Web Knowledge Graphs, describes a list of applications that exploit Semantic Web knowledge graphs like classification and regression, showing that the approaches developed in Part I and Part II can be used in applications in various domains. The book will be of interest to all those working in the field of data mining and KDD.

Probabilistic Databases Dan Suciu 2022-05-31
Probabilistic databases are databases where the

value of some attributes or the presence of some records are uncertain and known only with some probability. Applications in many areas such as information extraction, RFID and scientific data management, data cleaning, data integration, and financial risk assessment produce large volumes of uncertain data, which are best modeled and processed by a probabilistic database. This book presents the state of the art in representation formalisms and query processing techniques for probabilistic data. It starts by discussing the basic principles for representing large probabilistic databases, by decomposing them into tuple-independent tables, block-independent-disjoint tables, or U-databases. Then it discusses two classes of techniques for query evaluation on probabilistic databases. In extensional query evaluation, the entire probabilistic inference can be pushed into the database engine and, therefore, processed as effectively as the evaluation of standard SQL queries. The relational queries that can be

evaluated this way are called safe queries. In intensional query evaluation, the probabilistic inference is performed over a propositional formula called lineage expression: every relational query can be evaluated this way, but the data complexity dramatically depends on the query being evaluated, and can be #P-hard. The book also discusses some advanced topics in probabilistic data management such as top-k query processing, sequential probabilistic databases, indexing and materialized views, and Monte Carlo databases. Table of Contents: Overview / Data and Query Model / The Query Evaluation Problem / Extensional Query Evaluation / Intensional Query Evaluation / Advanced Techniques
Entity-Oriented Search - Krisztian Balog
2018-10-02

This open access book covers all facets of entity-oriented search—where “search” can be interpreted in the broadest sense of information access—from a unified point of view, and

provides a coherent and comprehensive overview of the state of the art. It represents the first synthesis of research in this broad and rapidly developing area. Selected topics are discussed in-depth, the goal being to establish fundamental techniques and methods as a basis for future research and development. Additional topics are treated at a survey level only, containing numerous pointers to the relevant literature. A roadmap for future research, based on open issues and challenges identified along the way, rounds out the book. The book is divided into three main parts, sandwiched between introductory and concluding chapters. The first two chapters introduce readers to the basic concepts, provide an overview of entity-oriented search tasks, and present the various types and sources of data that will be used throughout the book. Part I deals with the core task of entity ranking: given a textual query, possibly enriched with additional elements or structural hints, return a ranked list of entities.

This core task is examined in a number of different variants, using both structured and unstructured data collections, and numerous query formulations. In turn, Part II is devoted to the role of entities in bridging unstructured and structured data. Part III explores how entities can enable search engines to understand the concepts, meaning, and intent behind the query that the user enters into the search box, and how they can provide rich and focused responses (as opposed to merely a list of documents)—a process known as semantic search. The final chapter concludes the book by discussing the limitations of current approaches, and suggesting directions for future research. Researchers and graduate students are the primary target audience of this book. A general background in information retrieval is sufficient to follow the material, including an understanding of basic probability and statistics concepts as well as a basic knowledge of machine learning concepts and supervised

learning algorithms.

Graph Data Management - George Fletcher
2018-10-31

This book presents a comprehensive overview of fundamental issues and recent advances in graph data management. Its aim is to provide beginning researchers in the area of graph data management, or in fields that require graph data management, an overview of the latest developments in this area, both in applied and in fundamental subdomains. The topics covered range from a general introduction to graph data management, to more specialized topics like graph visualization, flexible queries of graph data, parallel processing, and benchmarking. The book will help researchers put their work in perspective and show them which types of tools, techniques and technologies are available, which ones could best suit their needs, and where there are still open issues and future research directions. The chapters are contributed by leading experts in the relevant areas, presenting

a coherent overview of the state of the art in the field. Readers should have a basic knowledge of data management techniques as they are taught in computer science MSc programs.

Database Systems: The Complete Book -

Hector Garcia-Molina 2008

The Past Web - Daniel Gomes 2021-06-30

This book provides practical information about web archives, offers inspiring examples for web archivists, raises new challenges, and shares recent research results about access methods to explore information from the past preserved by web archives. The book is structured in six parts. Part 1 advocates for the importance of web archives to preserve our collective memory in the digital era, demonstrates the problem of web ephemera and shows how web archiving activities have been trying to address this challenge. Part 2 then focuses on different strategies for selecting web content to be preserved and on the media types that different

web archives host. It provides an overview of efforts to address the preservation of web content as well as smaller-scale but high-quality collections of social media or audiovisual content. Next, Part 3 presents examples of initiatives to improve access to archived web information and provides an overview of access mechanisms for web archives designed to be used by humans or automatically accessed by machines. Part 4 presents research use cases for web archives. It also discusses how to engage more researchers in exploiting web archives and provides inspiring research studies performed using the exploration of web archives. Subsequently, Part 5 demonstrates that web archives should become crucial infrastructures for modern connected societies. It makes the case for developing web archives as research infrastructures and presents several inspiring examples of added-value services built on web archives. Lastly, Part 6 reflects on the evolution of the web and the sustainability of web

archiving activities. It debates the requirements and challenges for web archives if they are to assume the responsibility of being societal infrastructures that enable the preservation of memory. This book targets academics and advanced professionals in a broad range of research areas such as digital humanities, social sciences, history, media studies and information or computer science. It also aims to fill the need for a scholarly overview to support lecturers who would like to introduce web archiving into their courses by offering an initial reference for students.

Hadoop in Practice Alex Holmes 2014-09-29 Summary Hadoop in Practice, Second Edition provides over 100 tested, instantly useful techniques that will help you conquer big data, using Hadoop. This revised new edition covers changes and new features in the Hadoop core architecture, including MapReduce 2. Brand new chapters cover YARN and integrating Kafka, Impala, and Spark SQL with Hadoop. You'll also

get new and updated techniques for Flume, Sqoop, and Mahout, all of which have seen major new versions recently. In short, this is the most practical, up-to-date coverage of Hadoop available anywhere. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Book It's always a good time to upgrade your Hadoop skills! Hadoop in Practice, Second Edition provides a collection of 104 tested, instantly useful techniques for analyzing real-time streams, moving data securely, machine learning, managing large-scale clusters, and taming big data using Hadoop. This completely revised edition covers changes and new features in Hadoop core, including MapReduce 2 and YARN. You'll pick up hands-on best practices for integrating Spark, Kafka, and Impala with Hadoop, and get new and updated techniques for the latest versions of Flume, Sqoop, and Mahout. In short, this is the most practical, up-to-date coverage of Hadoop available. Readers

need to know a programming language like Java and have basic familiarity with Hadoop. What's Inside Thoroughly updated for Hadoop 2 How to write YARN applications Integrate real-time technologies like Storm, Impala, and Spark Predictive analytics using Mahout and RR Readers need to know a programming language like Java and have basic familiarity with Hadoop. About the Author Alex Holmes works on tough big-data problems. He is a software engineer, author, speaker, and blogger specializing in large-scale Hadoop projects. Table of Contents PART 1 BACKGROUND AND FUNDAMENTALS Hadoop in a heartbeat Introduction to YARN PART 2 DATA LOGISTICS Data serialization—working with text and beyond Organizing and optimizing data in HDFS Moving data into and out of Hadoop PART 3 BIG DATA PATTERNS Applying MapReduce patterns to big data Utilizing data structures and algorithms at scale Tuning, debugging, and testing PART 4 BEYOND MAPREDUCE SQL on Hadoop Writing

a YARN application Biologically Inspired Cognitive Architectures 2019 - Alexei V. Samsonovich 2019-07-16 The book focuses on original approaches intended to support the development of biologically inspired cognitive architectures. It bridges together different disciplines, from classical artificial intelligence to linguistics, from neuro- and social sciences to design and creativity, among others. The chapters, based on contributions presented at the Tenth Annual Meeting of the BICA Society, held in on August 15-18, 2019, in Seattle, WA, USA, discuss emerging methods, theories and ideas towards the realization of general-purpose humanlike artificial intelligence or fostering a better understanding of the ways the human mind works. All in all, the book provides engineers, mathematicians, psychologists, computer scientists and other experts with a timely snapshot of recent research and a source of inspiration for future developments in the

broadly intended areas of artificial intelligence and biological inspiration.

Handbook of Satisfiability - A. Biere

2021-05-05

Propositional logic has been recognized throughout the centuries as one of the cornerstones of reasoning in philosophy and mathematics. Over time, its formalization into Boolean algebra was accompanied by the recognition that a wide range of combinatorial problems can be expressed as propositional satisfiability (SAT) problems. Because of this dual role, SAT developed into a mature, multi-faceted scientific discipline, and from the earliest days of computing a search was underway to discover how to solve SAT problems in an automated fashion. This book, the Handbook of Satisfiability, is the second, updated and revised edition of the book first published in 2009 under the same name. The handbook aims to capture the full breadth and depth of SAT and to bring together significant

progress and advances in automated solving. Topics covered span practical and theoretical research on SAT and its applications and include search algorithms, heuristics, analysis of algorithms, hard instances, randomized formulae, problem encodings, industrial applications, solvers, simplifiers, tools, case studies and empirical results. SAT is interpreted in a broad sense, so as well as propositional satisfiability, there are chapters covering the domain of quantified Boolean formulae (QBF), constraints programming techniques (CSP) for word-level problems and their propositional encoding, and satisfiability modulo theories (SMT). An extensive bibliography completes each chapter. This second edition of the handbook will be of interest to researchers, graduate students, final-year undergraduates, and practitioners using or contributing to SAT, and will provide both an inspiration and a rich resource for their work. Edmund Clarke, 2007 ACM Turing Award Recipient: "SAT solving is a

key technology for 21st century computer science." Donald Knuth, 1974 ACM Turing Award Recipient: "SAT is evidently a killer app, because it is key to the solution of so many other problems." Stephen Cook, 1982 ACM Turing Award Recipient: "The SAT problem is at the core of arguably the most fundamental question in computer science: What makes a problem hard?"

Digital Forensic Evidence Examination -

Fred Cohen 2012-03-01

Digital Forensic Evidence Examination focuses on the scientific basis for analysis, interpretation, attribution, and reconstruction of digital forensic evidence in a legal context. It defines the bounds of "Information Physics" as it affects digital forensics, describes a model of the overall processes associated with the use of such evidence in legal matters, and provides the detailed basis for the science of digital forensic evidence examination. It reviews and discusses digital forensic evidence analysis, interpretation,

attribution, and reconstruction and their scientific bases, discusses tools and methodologies and their limits, and reviews the state of the science and its future outlook.

Mining the Web - Soumen Chakrabarti

2002-10-09

The definitive book on mining the Web from the preeminent authority.

Data Streams - Charu C. Aggarwal 2007-04-03

This book primarily discusses issues related to the mining aspects of data streams and it is unique in its primary focus on the subject. This volume covers mining aspects of data streams comprehensively: each contributed chapter contains a survey on the topic, the key ideas in the field for that particular topic, and future research directions. The book is intended for a professional audience composed of researchers and practitioners in industry. This book is also appropriate for advanced-level students in computer science.

Core Data Analysis: Summarization,

Correlation, and Visualization - Boris Mirkin
2019-04-15

This text examines the goals of data analysis with respect to enhancing knowledge, and identifies data summarization and correlation analysis as the core issues. Data summarization, both quantitative and categorical, is treated within the encoder-decoder paradigm bringing forward a number of mathematically supported insights into the methods and relations between them. Two Chapters describe methods for categorical summarization: partitioning, divisive clustering and separate cluster finding and another explain the methods for quantitative summarization, Principal Component Analysis and PageRank. Features: · An in-depth presentation of K-means partitioning including a corresponding Pythagorean decomposition of the data scatter. · Advice regarding such issues as clustering of categorical and mixed scale data, similarity and network data, interpretation aids, anomalous clusters, the number of clusters, etc.

· Thorough attention to data-driven modelling including a number of mathematically stated relations between statistical and geometrical concepts including those between goodness-of-fit criteria for decision trees and data standardization, similarity and consensus clustering, modularity clustering and uniform partitioning. New edition highlights: · Inclusion of ranking issues such as Google PageRank, linear stratification and tied rankings median, consensus clustering, semi-average clustering, one-cluster clustering · Restructured to make the logics more straightforward and sections self-contained Core Data Analysis:

Summarization, Correlation and Visualization is aimed at those who are eager to participate in developing the field as well as appealing to novices and practitioners.

Legal Informatics - Daniel Martin Katz
2021-02-18

This cutting-edge volume offers a theoretical and applied introduction to the emerging legal

technology and informatics industry.

Knowledge Graphs - Aidan Hogan 2021-11-08

This book provides a comprehensive and accessible introduction to knowledge graphs, which have recently garnered notable attention from both industry and academia. Knowledge graphs are founded on the principle of applying a graph-based abstraction to data, and are now broadly deployed in scenarios that require integrating and extracting value from multiple, diverse sources of data at large scale. The book defines knowledge graphs and provides a high-level overview of how they are used. It presents and contrasts popular graph models that are commonly used to represent data as graphs, and the languages by which they can be queried before describing how the resulting data graph can be enhanced with notions of schema, identity, and context. The book discusses how ontologies and rules can be used to encode knowledge as well as how inductive techniques—based on statistics, graph analytics,

machine learning, etc.—can be used to encode and extract knowledge. It covers techniques for the creation, enrichment, assessment, and refinement of knowledge graphs and surveys recent open and enterprise knowledge graphs and the industries or applications within which they have been most widely adopted. The book closes by discussing the current limitations and future directions along which knowledge graphs are likely to evolve. This book is aimed at students, researchers, and practitioners who wish to learn more about knowledge graphs and how they facilitate extracting value from diverse data at large scale. To make the book accessible for newcomers, running examples and graphical notation are used throughout. Formal definitions and extensive references are also provided for those who opt to delve more deeply into specific topics.

Introduction to Pattern Recognition and Machine Learning - M Narasimha Murty

2015-04-22

Downloaded from constructivworks.com
on by guest

This book adopts a detailed and methodological algorithmic approach to explain the concepts of pattern recognition. While the text provides a systematic account of its major topics such as pattern representation and nearest neighbour based classifiers, current topics — neural networks, support vector machines and decision trees — attributed to the recent vast progress in this field are also dealt with. Introduction to Pattern Recognition and Machine Learning will equip readers, especially senior computer science undergraduates, with a deeper understanding of the subject matter.

Contents: Introduction Types of Data Feature Extraction and Feature Selection Bayesian Learning Classification Classification Using Soft Computing Techniques Data Clustering Soft Clustering Application — Social and Information Networks Readership: Academics and working professionals in computer science. Key Features: The algorithmic approach taken and the practical issues dealt with will aid the reader

in writing programs and implementing methods Covers recent and advanced topics by providing working exercises, examples and illustrations in each chapter Provides the reader with a deeper understanding of the subject matter

Keywords: Clustering; Classification; Supervised Learning; Soft Computing

Handbuch Hochschullehre Digital - Jürgen Handke 2020-05-08

Schon seit Jahren ist die Digitalisierung der Lehre ein wichtiges Thema an den Hochschulen, hatte sich in der Anwendung aber noch nicht richtig durchgesetzt. Getrieben durch die Herausforderungen der Corona-Krise sind Ideen zur Umsetzung plötzlich dringend notwendig. Mit dieser Anleitung zeigt Prof. Dr. Jürgen Handke Wege auf, wie der Einstieg in die Digitalisierung gelingen kann. Ausgehend von Problemen der Hochschullehre diskutiert er die Möglichkeiten anhand von konkreten Beispielen. Es gilt, mit klugen Digitalisierungskonzepten die Probleme der klassischen Lehre zu überwinden,

sodass schließlich im Zusammenspiel von bewährten analogen Konzepten und digitalen Inhaltsvermittlungsformaten signifikante Mehrwerte entstehen.

The Foundations for Provenance on the Web

- Luc Moreau 2010-08-26

Aimed at anyone who discovers or publishes information on the Web and who cares about its origin and its quality. Based on an analysis of literature, this survey puts forward the Open Provenance Vision.

ICT in Education, Research, and Industrial Applications - Vadim Ermolayev 2013-01-11

This book constitutes the refereed proceedings of the 8th International Conference on ICT in

Education, Research, and Industrial Applications, held in Kherson, Ukraine, in June 2012. The 14 revised full papers were carefully reviewed and selected from 70 submissions. This book begins with an invited contribution presenting the substance of one of ICTERI 2012 invited talks. The chapter deals with the issues of abstraction and verification of properties in real-time Java programs. The rest of the volume is structured in four topical parts: ICT Frameworks, Infrastructures, Integration, and Deployment; Formal Logic and Knowledge-Based Frameworks; ICT-Based Systems Modeling, Specification, and Verification; ICT in Teaching and Learning.