14 Predicate Logic Stanford Infolab

Pervasive Computing
Cooperative Information Systems
Exploiting Semantic Web Knowledge Graphs in Data Mining
Hadoop in Practice
Handbook of Satisfiability
Third International Conference, PERVASIVE 2005, Munich, Germany, May 8-13, 2005, Proceedings
Making Institutional Repositories Work
Readings in Database Systems
Big Data 2.0 Processing Systems
Results of the LOD2 Project
A Semantic Web Primer, third edition
Metadata and Semantic Research
Graph Data Management
A Mechanology of Algorithmic Techniques
Leitfaden für eine moderne und mediengerechte Lehre
Learning Action Models for Reactive Autonomous Agents

ALGORITHMS OF THE INTELLIGENT WEB
Foundations of Fuzzy Logic and Semantic Web Languages (Open Access)
Third International Workshop, CIA'99 Uppsala, Sweden, July 31 - August 2, 1999 Proceedings
8th International Conference, ICTERI 2012, Kherson, Ukraine, June 6-10, 2012, Revised Selected Papers
Introduction to Pattern Recognition and Machine Learning
Context and Interest in a Grammar of Multimodal Meaning
Handbuch Hochschullehre Digital
Foundations of Computer Science
Semantics-Oriented Natural Language Processing
Probabilistic Databases
9th International Conference on NLP, PolTAL 2014, Warsaw, Poland, September 17-19, 2014. Proceedings
Entity-Oriented Search
Mathematical Models and Algorithms
C Edition
Fundamental Issues and Recent Developments
Exploring Web Archives
12th International Conference on Conceptual Structures, ICCS 2004, Huntsville, AL,

USA, July 19-23, 2004, Proceedings
Legal Informatics
9th International Conference, CoopIS 2001, Trento, Italy, September 5-7, 2001. Proceedings
The Beauty of Mathematics in Computer Science
Linked Open Data -- Creating Knowledge Out of Interlinked Data
14th International Conference, MTSR 2020, Madrid, Spain, December 2–4, 2020, Revised Selected Papers
Adding Sense

RILEY SCHWARTZ

Pervasive Computing Fred Cohen & Associates
This book constitutes the refereed proceedings of the Third International Workshop on Cooperative Information Systems, CIA'99, held in Uppsala, Sweden in July/August 1999. The 16 revised full papers presented were carefully reviewed and selected from a total of 46 submissions. Also included are ten invited contributions by leading experts. The volume is divided in sections on information discovery and management on the Internet; information agents on the Internet-prototypes systems and applications;

communication and collaboration, mobile information agents; rational information agents for electronic business; service mediation and negotiation; and adaptive personal assistance.

Cooperative Information Systems Cambridge University Press

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.

Exploiting Semantic Web Knowledge Graphs in Data Mining World Scientific

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.

Hadoop in Practice Purdue University Press

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.

Handbook of Satisfiability Morgan & Claypool Publishers

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.
Third International Conference, PERVASIVE 2005, Munich, Germany, May 8-13, 2005, Proceedings Morgan & Claypool Publishers
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.

Making Institutional Repositories Work Pearson Higher Ed

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.
Readings in Database Systems W. H. Freeman
The Beauty of Mathematics in Computer Science explains the mathematical fundamentals of information technology products and services we use every day, from Google Web Search to GPS Navigation, and from speech recognition to CDMA mobile services. The book was published in Chinese in 2011 and has sold more than 600,000 copies. Readers were surprised to find that many daily-used IT technologies were so tightly tied to mathematical principles. For example, the automatic classification of news articles uses the cosine law taught in high school. The book covers many topics related to computer applications and applied mathematics including: Natural language processing Speech recognition and machine translation Statistical language modeling Quantitive measurement of information Graph theory and web crawler Pagerank for web search Matrix operation and document classification Mathematical background of big data Neural networks and Google's deep learning Jun Wu was a staff research scientist in Google who invented Google's Chinese, Japanese, and Korean Web Search Algorithms and was responsible for many Google

machine learning projects. He wrote official blogs introducing Google technologies behind its products in very simple languages for Chinese Internet users from 2006-2010. The blogs had more than 2 million followers. Wu received PhD in computer science from Johns Hopkins University and has been working on speech recognition and natural language processing for more than 20 years. He was one of the earliest engineers of Google, managed many products of the company, and was awarded 19 US patents during his 10-year tenure there. Wu became a full-time VC investor and co-founded Amino Capital in Palo Alto in 2014 and is the author of eight books.

Big Data 2.0 Processing Systems Linköping University Electronic Press

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

Results of the LOD2 Project Tectum Wissenschaftsverlag

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.

A Semantic Web Primer, third edition

Springer
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

Metadata and Semantic Research

Handbuch Hochschullehre DigitalLeitfaden für eine moderne und mediengerechte Lehre

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:IntroductionTypes of DataFeature Extraction and Feature SelectionBayesian LearningClassificationClassification Using Soft Computing TechniquesData ClusteringSoft ClusteringApplication — 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 methodsCovers recent and advanced topics by providing working exercises, examples and illustrations in each chapterProvides the reader with a deeper understanding of the subject matterKeywords:Clustering;Classification ;Supervised Learning;Soft Computing

Graph Data Management IOS Press

The Semantic Web is characterized by the existence of a very large number of distributed semantic resources, which together define a network of ontologies. These ontologies in turn are interlinked through a variety of different meta-relationships such as versioning, inclusion, and many more. This scenario is radically different from the relatively narrow contexts in which ontologies have been traditionally developed and applied, and thus calls for new methods and tools to effectively support the development of novel network-oriented

semantic applications. This book by Suárez-Figueroa et al. provides the necessary methodological and technological support for the development and use of ontology networks, which ontology developers need in this distributed environment. After an introduction, in its second part the authors describe the NeOn Methodology framework. The book's third part details the key activities relevant to the ontology engineering life cycle. For each activity, a general introduction, methodological guidelines, and practical examples are provided. The fourth part then presents a detailed overview of the NeOn Toolkit and its plug-ins. Lastly, case studies from the pharmaceutical and the fishery domain round out the work. The book primarily addresses two main audiences: students (and their lecturers) who need a textbook for advanced undergraduate or graduate courses on ontology engineering, and practitioners who need to develop ontologies in particular or Semantic Web-based applications in general. Its educational value is maximized by its structured approach to explaining guidelines and combining them with case studies and numerous examples. The description of the open source NeOn Toolkit provides an additional asset, as it allows readers to easily evaluate and apply the ideas presented.

Springer Science & Business Media

This is the eBook of the printed book and may not include any media, website access codes, or print supplements that

may come packaged with the bound book. Database Systems: The Complete Book is ideal for Database Systems and Database Design and Application courses offered at the junior, senior and graduate levels in Computer Science departments. A basic understanding of algebraic expressions and laws, logic, basic data structure, OOP concepts, and programming environments is implied. Written by well-known computer scientists, this introduction to database systems offers a comprehensive approach, focusing on database design, database use, and implementation of database applications and database management systems. The first half of the book provides in-depth coverage of databases from the point of view of the database designer, user, and application programmer. It covers the latest database standards SQL:1999, SQL/PSM, SQL/CLI, JDBC, ODL, and XML, with broader coverage of SQL than most other texts. The second half of the book provides in-depth coverage of databases from the point of view of the DBMS implementor. It focuses on storage structures, query processing, and transaction management. The book covers the main techniques in these areas with broader coverage of query optimization than most other texts, along with advanced topics including multidimensional and bitmap indexes, distributed transactions, and information integration techniques.

A Mechanology of Algorithmic Techniques Springer Science & Business Media

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.

Leitfaden für eine moderne und mediengerechte Lehre Springer Science & Business Media

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.

Learning Action Models for Reactive Autonomous Agents Springer Nature

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

ALGORITHMS OF THE INTELLIGENT WEB
MIT Press

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.

Foundations of Fuzzy Logic and Semantic Web Languages (Open Access) Springer Science & Business Media

Previous work on action-model learning has focused on domains that contain only deterministic, atomic action models that explicitly describe all changes that can occur in the environment. The thesis extends this previous work to cover domains that contain durative actions, continuous variables, nondeterministic action effects, and actions taken by other agents. Results have been demonstrated in several robot simulation environments and the Silicon Graphics, Inc. flight simulator.

Third International Workshop, CIA'99 Uppsala, Sweden, July 31 - August 2, 1999 Proceedings CRC Press

This groundbreaking work offers a first-of-its-kind overview of legal informatics, the academic discipline underlying the

technological transformation and economics of the legal industry. Edited by Daniel Martin Katz, Ron Dolin, and Michael J. Bommarito, and featuring contributions from more than two dozen academic and industry experts, chapters cover the history and principles of legal informatics and background technical concepts – including natural language processing and distributed ledger technology. The volume also presents real-world case studies that offer important insights into document review, due diligence, compliance, case prediction, billing, negotiation and settlement, contracting, patent management, legal research, and online dispute resolution. Written for both technical and non-technical readers, Legal Informatics is the ideal resource for anyone interested in identifying, understanding, and executing opportunities in this exciting field.