

Learning To Rank For Information Retrieval And Natural Language Processing Second Edition Synthesis Lectures On Human Language Technologies

Right here, we have countless book **Learning To Rank For Information Retrieval And Natural Language Processing Second Edition Synthesis Lectures On Human Language Technologies** and collections to check out. We additionally provide variant types and afterward type of the books to browse. The usual book, fiction, history, novel, scientific research, as well as various supplementary sorts of books are readily understandable here.

As this Learning To Rank For Information Retrieval And Natural Language Processing Second Edition Synthesis Lectures On Human Language Technologies , it ends happening brute one of the favored book Learning To Rank For Information Retrieval And Natural Language Processing Second Edition Synthesis Lectures On Human Language Technologies collections that we have. This is why you remain in the best website to see the unbelievable ebook to have.

Learning to Rank for Information Retrieval - Tie-Yan Liu 2011-04-29

Due to the fast growth of the Web and the difficulties in finding desired information, efficient and effective information retrieval systems have become more important than ever, and the search engine has become an essential tool for many people. The ranker, a central component in every search engine, is responsible for the matching between processed queries and indexed documents. Because of its central role, great attention has been paid to the research and development of ranking technologies. In addition, ranking is also pivotal for many other information retrieval applications, such as collaborative filtering, definition ranking, question answering, multimedia retrieval, text summarization, and online advertisement. Leveraging machine learning technologies in the ranking process has led to innovative and more effective ranking models, and eventually to a completely new research area called "learning to rank". Liu first gives a comprehensive review of the major approaches to learning to rank. For each approach he presents the basic framework, with example algorithms, and he discusses its advantages and disadvantages. He continues with some recent advances in learning to rank that cannot be simply categorized into the three major approaches - these include relational ranking, query-dependent ranking, transfer ranking, and semisupervised ranking. His presentation is completed by several examples that apply these technologies to solve real information retrieval problems, and by theoretical discussions on guarantees for ranking performance. This book is written for researchers and graduate students in both information retrieval and machine learning. They will find here the only comprehensive description of the state of the art in a field that has driven the recent advances in search engine development.

Knowledge Graphs - Ridho Reinanda 2020-10-15

The authors of this survey present an extensive overview of tasks related to KGs from an IR perspective, provide a thorough review for each task, and present discussions on common issues that are shared among the tasks.

Python for Everybody - Charles R. Severance 2016-04-09

Python for Everybody is designed to introduce students to programming and software development through the lens of exploring data. You can think of the Python programming language as your tool to solve data problems that are beyond the capability of a spreadsheet. Python is an easy to use and easy to learn programming language that is freely available on Macintosh, Windows, or Linux computers. So once you learn Python you can use it for the rest of your career without needing to purchase any software. This book uses the Python 3 language. The earlier Python 2 version of this book is titled "Python for Informatics: Exploring Information". There are free downloadable electronic copies of this book in various formats and supporting materials for the book at www.pythonlearn.com. The course materials are available to you under a Creative Commons License so you can adapt them to teach your own Python course.

Aggregated Search - Jaime Arguello 2017-03-06

A comprehensive summary of the research in aggregated search. It also surveys different evaluation methodologies for aggregated search and discusses prior user studies aimed at better understanding how users behave with aggregated search interfaces. It concludes by reviewing different advanced topics in aggregated search.

Deep Learning for Search - Tommaso Teofili 2019-06-02

Summary Deep Learning for Search teaches you how to improve the

effectiveness of your search by implementing neural network-based techniques. By the time you're finished with the book, you'll be ready to build amazing search engines that deliver the results your users need and that get better as time goes on! Foreword by Chris Mattmann. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Deep learning handles the toughest search challenges, including imprecise search terms, badly indexed data, and retrieving images with minimal metadata. And with modern tools like DL4J and TensorFlow, you can apply powerful DL techniques without a deep background in data science or natural language processing (NLP). This book will show you how. About the Book Deep Learning for Search teaches you to improve your search results with neural networks. You'll review how DL relates to search basics like indexing and ranking. Then, you'll walk through in-depth examples to upgrade your search with DL techniques using Apache Lucene and Deeplearning4j. As the book progresses, you'll explore advanced topics like searching through images, translating user queries, and designing search engines that improve as they learn! What's inside Accurate and relevant rankings Searching across languages Content-based image search Search with recommendations About the Reader For developers comfortable with Java or a similar language and search basics. No experience with deep learning or NLP needed. About the Author Tommaso Teofili is a software engineer with a passion for open source and machine learning. As a member of the Apache Software Foundation, he contributes to a number of open source projects, ranging from topics like information retrieval (such as Lucene and Solr) to natural language processing and machine translation (including OpenNLP, Joshua, and UIMA). He currently works at Adobe, developing search and indexing infrastructure components, and researching the areas of natural language processing, information retrieval, and deep learning. He has presented search and machine learning talks at conferences including BerlinBuzzwords, International Conference on Computational Science, ApacheCon, EclipseCon, and others. You can find him on Twitter at @tteofili. Table of Contents PART 1 - SEARCH MEETS DEEP LEARNING Neural search Generating synonyms PART 2 - THROWING NEURAL NETS AT A SEARCH ENGINE From plain retrieval to text generation More-sensitive query suggestions Ranking search results with word embeddings Document embeddings for rankings and recommendations PART 3 - ONE STEP BEYOND Searching across languages Content-based image search A peek at performance **Advances in Knowledge Discovery and Data Mining** - Thanaruk Theeramunkong 2009-04-20 The Pacific-Asia Conference on Knowledge Discovery and Data Mining has been held every year from 1997. PAKDD 2009, the 13th in the series, was held in Bangkok, Thailand during April 27-30, 2008. PAKDD is a major international conference in the areas of data mining (DM) and knowledge discovery in database (KDD). It provides an international forum for researchers and industry practitioners to share their new ideas, original research results and practical development experiences from all KDD-related areas including data mining, data warehousing, machine learning, databases, statistics, knowledge acquisition and automatic scientific discovery, data visualization, causal induction and knowledge-based systems. For PAKDD 2009, we received 338 research papers from various countries and regions in Asia, Australia, North America, South America, Europe, and Africa. Every submission was rigorously reviewed by at least three

reviewers with a double-blind protocol. The initial results were discussed among the reviewers and finally judged by the Program Committee Chairs. When there was a conflict, an additional review was provided by the Program Committee Chairs. The Program Committee members were deeply involved in the highly selective process. As a result, only 39 papers (approximately 11.5% of the 338 submitted papers) were accepted as regular papers, 73 papers (21.6% of them) were accepted as short papers.

2019 Amity International Conference on Artificial Intelligence (AICAI). - 2019

Learning to Rank for Information Retrieval and Natural Language Processing - Hang Li 2011-04-20

Learning to rank refers to machine learning techniques for training the model in a ranking task. Learning to rank is useful for many applications in information retrieval, natural language processing, and data mining. Intensive studies have been conducted on the problem recently and significant progress has been made. This lecture gives an introduction to the area including the fundamental problems, existing approaches, theories, applications, and future work. The author begins by showing that various ranking problems in information retrieval and natural language processing can be formalized as two basic ranking tasks, namely ranking creation (or simply ranking) and ranking aggregation. In ranking creation, given a request, one wants to generate a ranking list of offerings based on the features derived from the request and the offerings. In ranking aggregation, given a request, as well as a number of ranking lists of offerings, one wants to generate a new ranking list of the offerings. Ranking creation (or ranking) is the major problem in learning to rank. It is usually formalized as a supervised learning task. The author gives detailed explanations on learning for ranking creation and ranking aggregation, including training and testing, evaluation, feature creation, and major approaches. Many methods have been proposed for ranking creation. The methods can be categorized as the pointwise, pairwise, and listwise approaches according to the loss functions they employ. They can also be categorized according to the techniques they employ, such as the SVM based, Boosting SVM, Neural Network based approaches. The author also introduces some popular learning to rank methods in details. These include PRank, OC SVM, Ranking SVM, IR SVM, GBRank, RankNet, LambdaRank, ListNet & ListMLE, AdaRank, SVM MAP, SoftRank, Borda Count, Markov Chain, and CRanking. The author explains several example applications of learning to rank including web search, collaborative filtering, definition search, keyphrase extraction, query dependent summarization, and re-ranking in machine translation. A formulation of learning for ranking creation is given in the statistical learning framework. Ongoing and future research directions for learning to rank are also discussed. Table of Contents: Introduction / Learning for Ranking Creation / Learning for Ranking Aggregation / Methods of Learning to Rank / Applications of Learning to Rank / Theory of Learning to Rank / Ongoing and Future Work

Relevance Ranking for Vertical Search Engines - Bo Long 2014-01-25

In plain, uncomplicated language, and using detailed examples to explain the key concepts, models, and algorithms in vertical search ranking, *Relevance Ranking for Vertical Search Engines* teaches readers how to manipulate ranking algorithms to achieve better results in real-world applications. This reference book for professionals covers concepts and theories from the fundamental to the advanced, such as relevance, query intention, location-based relevance ranking, and cross-property ranking. It covers the most recent developments in vertical search ranking applications, such as freshness-based relevance theory for new search applications, location-based relevance theory for local search applications, and cross-property ranking theory for applications involving multiple verticals. Foreword by Ron Brachman, Chief Scientist and Head, Yahoo! Labs Introduces ranking algorithms and teaches readers how to manipulate ranking algorithms for the best results Covers concepts and theories from the fundamental to the advanced Discusses the state of the art: development of theories and practices in vertical search ranking applications Includes detailed examples, case studies and real-world situations

Search Engines - Bruce Croft 2011-11-21

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. *Search Engines: Information Retrieval in Practice* is ideal for introductory information retrieval courses at the

undergraduate and graduate level in computer science, information science and computer engineering departments. It is also a valuable tool for search engine and information retrieval professionals. Written by a leader in the field of information retrieval, *Search Engines: Information Retrieval in Practice*, is designed to give undergraduate students the understanding and tools they need to evaluate, compare and modify search engines. Coverage of the underlying IR and mathematical models reinforce key concepts. The book's numerous programming exercises make extensive use of Galago, a Java-based open source search engine. **Click Models for Web Search** - Aleksandr Chuklin 2015-07-01 With the rapid growth of web search in recent years the problem of modeling its users has started to attract more and more attention of the information retrieval community. This has several motivations. By building a model of user behavior we are essentially developing a better understanding of a user, which ultimately helps us to deliver a better search experience. A model of user behavior can also be used as a predictive device for non-observed items such as document relevance, which makes it useful for improving search result ranking. Finally, in many situations experimenting with real users is just infeasible and hence user simulations based on accurate models play an essential role in understanding the implications of algorithmic changes to search engine results or presentation changes to the search engine result page. In this survey we summarize advances in modeling user click behavior on a web search engine result page. We present simple click models as well as more complex models aimed at capturing non-trivial user behavior patterns on modern search engine result pages. We discuss how these models compare to each other, what challenges they have, and what ways there are to address these challenges. We also study the problem of evaluating click models and discuss the main applications of click models.

Solving Least Squares Problems - Charles L. Lawson 1995-12-01

This Classic edition includes a new appendix which summarizes the major developments since the book was originally published in 1974. The additions are organized in short sections associated with each chapter. An additional 230 references have been added, bringing the bibliography to over 400 entries. Appendix C has been edited to reflect changes in the associated software package and software distribution method.

Scaling Up Machine Learning - Ron Bekkerman 2012

This integrated collection covers a range of parallelization platforms, concurrent programming frameworks and machine learning settings, with case studies.

Statistical Language Models for Information Retrieval - ChengXiang Zhai 2009

As online information grows dramatically, search engines such as Google are playing a more and more important role in our lives. Critical to all search engines is the problem of designing an effective retrieval model that can rank documents accurately for a given query. This has been a central research problem in information retrieval for several decades. In the past ten years, a new generation of retrieval models, often referred to as statistical language models, has been successfully applied to solve many different information retrieval problems. Compared with the traditional models such as the vector space model, these new models have a more sound statistical foundation and can leverage statistical estimation to optimize retrieval parameters. They can also be more easily adapted to model non-traditional and complex retrieval problems. Empirically, they tend to achieve comparable or better performance than a traditional model with less effort on parameter tuning. This book systematically reviews the large body of literature on applying statistical language models to information retrieval with an emphasis on the underlying principles, empirically effective language models, and language models developed for non-traditional retrieval tasks. All the relevant literature has been synthesized to make it easy for a reader to digest the research progress achieved so far and see the frontier of research in this area. The book also offers practitioners an informative introduction to a set of practically useful language models that can effectively solve a variety of retrieval problems. No prior knowledge about information retrieval is required, but some basic knowledge about probability and statistics would be useful for fully digesting all the details. Table of Contents: Introduction / Overview of Information Retrieval Models / Simple Query Likelihood Retrieval Model / Complex Query Likelihood Model / Probabilistic Distance Retrieval Model / Language Models for Special Retrieval Tasks / Language Models for Latent Topic Analysis / Conclusions

Search Result Diversification - Santos Rodrygo L T 2015-02-27

This primer reviews the published literature on search result

diversification. In particular, it discusses the motivations for diversifying the search results for an ambiguous query and provides a formal definition of the search result diversification problem. In addition, it describes the most successful approaches in the literature for producing and evaluating diversity in multiple search domains.

Learning to Rank for Information Retrieval - Tie-Yan Liu 2009

Learning to Rank for Information Retrieval is an introduction to the field of learning to rank, a hot research topic in information retrieval and machine learning. It categorizes the state-of-the-art learning-to-rank algorithms into three approaches from a unified machine learning perspective, describes the loss functions and learning mechanisms in different approaches, reveals their relationships and differences, shows their empirical performances on real IR applications, and discusses their theoretical properties such as generalization ability. As a tutorial, Learning to Rank for Information Retrieval helps people find the answers to the following critical questions: To what respect are learning-to-rank algorithms similar and in which aspects do they differ? What are the strengths and weaknesses of each algorithm? Which learning-to-rank algorithm empirically performs the best? Is ranking a new machine learning problem? What are the unique theoretical issues for ranking as compared to classification and regression? Learning to Rank for Information Retrieval is both a guide for beginners who are embarking on research in this area, and a useful reference for established researchers and practitioners

Online Evaluation for Information Retrieval - Katja Hofmann
2016-06-07

Provides a comprehensive overview of the topic. It shows how online evaluation is used for controlled experiments, segmenting them into experiment designs that allow absolute or relative quality assessments. It also includes an extensive discussion of recent work on data re-use, and experiment estimation based on historical data.

Proceedings - Zoubin Ghahramani 2007

Advances in Information Retrieval - Joemon M. Jose 2020-04-10

This two-volume set LNCS 12035 and 12036 constitutes the refereed proceedings of the 42nd European Conference on IR Research, ECIR 2020, held in Lisbon, Portugal, in April 2020.* The 55 full papers presented together with 8 reproducibility papers, 46 short papers, 10 demonstration papers, 12 invited CLEF papers, 7 doctoral consortium papers, 4 workshop papers, and 3 tutorials were carefully reviewed and selected from 457 submissions. They were organized in topical sections named: Part I: deep learning I; entities; evaluation; recommendation; information extraction; deep learning II; retrieval; multimedia; deep learning III; queries; IR – general; question answering, prediction, and bias; and deep learning IV. Part II: reproducibility papers; short papers; demonstration papers; CLEF organizers lab track; doctoral consortium papers; workshops; and tutorials. *Due to the COVID-19 pandemic, this conference was held virtually.

Information Retrieval - Stefan Buttcher 2016-02-12

An introduction to information retrieval, the foundation for modern search engines, that emphasizes implementation and experimentation. Information retrieval is the foundation for modern search engines. This textbook offers an introduction to the core topics underlying modern search technologies, including algorithms, data structures, indexing, retrieval, and evaluation. The emphasis is on implementation and experimentation; each chapter includes exercises and suggestions for student projects. Wumpus—a multiuser open-source information retrieval system developed by one of the authors and available online—provides model implementations and a basis for student work. The modular structure of the book allows instructors to use it in a variety of graduate-level courses, including courses taught from a database systems perspective, traditional information retrieval courses with a focus on IR theory, and courses covering the basics of Web retrieval. In addition to its classroom use, Information Retrieval will be a valuable reference for professionals in computer science, computer engineering, and software engineering.

Learning to Rank for Information Retrieval and Natural Language Processing, Second Edition - Hang Li 2022-05-31

Learning to rank refers to machine learning techniques for training a model in a ranking task. Learning to rank is useful for many applications in information retrieval, natural language processing, and data mining. Intensive studies have been conducted on its problems recently, and significant progress has been made. This lecture gives an introduction to the area including the fundamental problems, major approaches, theories, applications, and future work. The author begins by showing

that various ranking problems in information retrieval and natural language processing can be formalized as two basic ranking tasks, namely ranking creation (or simply ranking) and ranking aggregation. In ranking creation, given a request, one wants to generate a ranking list of offerings based on the features derived from the request and the offerings. In ranking aggregation, given a request, as well as a number of ranking lists of offerings, one wants to generate a new ranking list of the offerings. Ranking creation (or ranking) is the major problem in learning to rank. It is usually formalized as a supervised learning task. The author gives detailed explanations on learning for ranking creation and ranking aggregation, including training and testing, evaluation, feature creation, and major approaches. Many methods have been proposed for ranking creation. The methods can be categorized as the pointwise, pairwise, and listwise approaches according to the loss functions they employ. They can also be categorized according to the techniques they employ, such as the SVM based, Boosting based, and Neural Network based approaches. The author also introduces some popular learning to rank methods in details. These include: PRank, OC SVM, McRank, Ranking SVM, IR SVM, GBRank, RankNet, ListNet & ListMLE, AdaRank, SVM MAP, SoftRank, LambdaRank, LambdaMART, Borda Count, Markov Chain, and CRanking. The author explains several example applications of learning to rank including web search, collaborative filtering, definition search, keyphrase extraction, query dependent summarization, and re-ranking in machine translation. A formulation of learning for ranking creation is given in the statistical learning framework. Ongoing and future research directions for learning to rank are also discussed. Table of Contents:

Learning to Rank / Learning for Ranking Creation / Learning for Ranking Aggregation / Methods of Learning to Rank / Applications of Learning to Rank / Theory of Learning to Rank / Ongoing and Future Work

Introduction to Information Retrieval - Christopher D. Manning
2008-07-07

Class-tested and coherent, this textbook teaches classical and web information retrieval, including web search and the related areas of text classification and text clustering from basic concepts. It gives an up-to-date treatment of all aspects of the design and implementation of systems for gathering, indexing, and searching documents; methods for evaluating systems; and an introduction to the use of machine learning methods on text collections. All the important ideas are explained using examples and figures, making it perfect for introductory courses in information retrieval for advanced undergraduates and graduate students in computer science. Based on feedback from extensive classroom experience, the book has been carefully structured in order to make teaching more natural and effective. Slides and additional exercises (with solutions for lecturers) are also available through the book's supporting website to help course instructors prepare their lectures.

Modern Information Retrieval - Yates 1999-09

Advanced Computing, Networking and Informatics- Volume 2 - Malay Kumar Kundu 2014-05-26

Advanced Computing, Networking and Informatics are three distinct and mutually exclusive disciplines of knowledge with no apparent sharing/overlap among them. However, their convergence is observed in many real world applications, including cyber-security, internet banking, healthcare, sensor networks, cognitive radio, pervasive computing amidst many others. This two-volume proceedings explore the combined use of Advanced Computing and Informatics in the next generation wireless networks and security, signal and image processing, ontology and human-computer interfaces (HCI). The two volumes together include 148 scholarly papers, which have been accepted for presentation from over 640 submissions in the second International Conference on Advanced Computing, Networking and Informatics, 2014, held in Kolkata, India during June 24-26, 2014. The first volume includes innovative computing techniques and relevant research results in informatics with selective applications in pattern recognition, signal/image processing and HCI. The second volume on the other hand demonstrates the possible scope of the computing techniques and informatics in wireless communications, networking and security.

Learning to Rank for Information Retrieval - Tie-Yan Liu 2011-05-06

Due to the fast growth of the Web and the difficulties in finding desired information, efficient and effective information retrieval systems have become more important than ever, and the search engine has become an essential tool for many people. The ranker, a central component in every search engine, is responsible for the matching between processed queries and indexed documents. Because of its central role, great

attention has been paid to the research and development of ranking technologies. In addition, ranking is also pivotal for many other information retrieval applications, such as collaborative filtering, definition ranking, question answering, multimedia retrieval, text summarization, and online advertisement. Leveraging machine learning technologies in the ranking process has led to innovative and more effective ranking models, and eventually to a completely new research area called "learning to rank". Liu first gives a comprehensive review of the major approaches to learning to rank. For each approach he presents the basic framework, with example algorithms, and he discusses its advantages and disadvantages. He continues with some recent advances in learning to rank that cannot be simply categorized into the three major approaches - these include relational ranking, query-dependent ranking, transfer ranking, and semisupervised ranking. His presentation is completed by several examples that apply these technologies to solve real information retrieval problems, and by theoretical discussions on guarantees for ranking performance. This book is written for researchers and graduate students in both information retrieval and machine learning. They will find here the only comprehensive description of the state of the art in a field that has driven the recent advances in search engine development.

Machine Learning: ECML 2007 - Joost N. Kok 2007-09-08

This book constitutes the refereed proceedings of the 18th European Conference on Machine Learning, ECML 2007, held in Warsaw, Poland, September 2007, jointly with PKDD 2007. The 41 revised full papers and 37 revised short papers presented together with abstracts of four invited talks were carefully reviewed and selected from 592 abstracts submitted to both, ECML and PKDD. The papers present a wealth of new results in the area and address all current issues in machine learning.

Genetic Programming IV - John R. Koza 2006-03-04

Genetic Programming IV: Routine Human-Competitive Machine Intelligence presents the application of GP to a wide variety of problems involving automated synthesis of controllers, circuits, antennas, genetic networks, and metabolic pathways. The book describes fifteen instances where GP has created an entity that either infringes or duplicates the functionality of a previously patented 20th-century invention, six instances where it has done the same with respect to post-2000 patented inventions, two instances where GP has created a patentable new invention, and thirteen other human-competitive results. The book additionally establishes: GP now delivers routine human-competitive machine intelligence GP is an automated invention machine GP can create general solutions to problems in the form of parameterized topologies GP has delivered qualitatively more substantial results in synchrony with the relentless iteration of Moore's Law

Information Retrieval with Verbose Queries - Manish Gupta 2015-07-31

The first monograph to provide a coherent and organized survey on this topic. It puts together the various research pieces of the puzzle, provides a comprehensive and structured overview of diverse proposed methods, and lists several application scenarios where effective verbose query processing can make a significant difference.

Occupational Outlook Handbook - United States. Bureau of Labor Statistics 1976

Machine Learning for Text - Charu C. Aggarwal 2018-03-19

Text analytics is a field that lies on the interface of information retrieval, machine learning, and natural language processing, and this textbook carefully covers a coherently organized framework drawn from these intersecting topics. The chapters of this textbook is organized into three categories: - Basic algorithms: Chapters 1 through 7 discuss the classical algorithms for machine learning from text such as preprocessing, similarity computation, topic modeling, matrix factorization, clustering, classification, regression, and ensemble analysis. - Domain-sensitive mining: Chapters 8 and 9 discuss the learning methods from text when combined with different domains such as multimedia and the Web. The problem of information retrieval and Web search is also discussed in the context of its relationship with ranking and machine learning methods. - Sequence-centric mining: Chapters 10 through 14 discuss various sequence-centric and natural language applications, such as feature engineering, neural language models, deep learning, text summarization, information extraction, opinion mining, text segmentation, and event detection. This textbook covers machine learning topics for text in detail. Since the coverage is extensive, multiple courses can be offered from the same book, depending on course level. Even though the presentation is text-centric,

Chapters 3 to 7 cover machine learning algorithms that are often used in domains beyond text data. Therefore, the book can be used to offer courses not just in text analytics but also from the broader perspective of machine learning (with text as a backdrop). This textbook targets graduate students in computer science, as well as researchers, professors, and industrial practitioners working in these related fields. This textbook is accompanied with a solution manual for classroom teaching.

Preference Learning - Johannes Fürnkranz 2010-11-19

The topic of preferences is a new branch of machine learning and data mining, and it has attracted considerable attention in artificial intelligence research in previous years. It involves learning from observations that reveal information about the preferences of an individual or a class of individuals. Representing and processing knowledge in terms of preferences is appealing as it allows one to specify desires in a declarative way, to combine qualitative and quantitative modes of reasoning, and to deal with inconsistencies and exceptions in a flexible manner. And, generalizing beyond training data, models thus learned may be used for preference prediction. This is the first book dedicated to this topic, and the treatment is comprehensive. The editors first offer a thorough introduction, including a systematic categorization according to learning task and learning technique, along with a unified notation. The first half of the book is organized into parts on label ranking, instance ranking, and object ranking; while the second half is organized into parts on applications of preference learning in multiattribute domains, information retrieval, and recommender systems. The book will be of interest to researchers and practitioners in artificial intelligence, in particular machine learning and data mining, and in fields such as multicriteria decision-making and operations research.

Advances in Information Retrieval - Djoerd Hiemstra 2021-03-26

This two-volume set LNCS 12656 and 12657 constitutes the refereed proceedings of the 43rd European Conference on IR Research, ECIR 2021, held virtually in March/April 2021, due to the COVID-19 pandemic. The 50 full papers presented together with 11 reproducibility papers, 39 short papers, 15 demonstration papers, 12 CLEF lab descriptions papers, 5 doctoral consortium papers, 5 workshop abstracts, and 8 tutorials abstracts were carefully reviewed and selected from 436 submissions. The accepted contributions cover the state of the art in IR: deep learning-based information retrieval techniques, use of entities and knowledge graphs, recommender systems, retrieval methods, information extraction, question answering, topic and prediction models, multimedia retrieval, and much more.

Modelling and Development of Intelligent Systems - Dana Simian 2020-01-16

This volume constitutes the refereed proceedings of the 6th International Conference on Modelling and Development of Intelligent Systems, MDIS 2019, held in Sibiu, Romania, in October 2019. The 13 revised full papers presented in the volume were carefully reviewed and selected from 31 submissions. The papers are organized in topical sections on adaptive systems; conceptual modelling; data mining; intelligent systems for decision support; machine learning.

Bandit Algorithms in Information Retrieval - Dorota Glowacka 2019-05-23

This monograph provides an overview of bandit algorithms inspired by various aspects of Information Retrieval. It is accessible to anyone who has completed introductory to intermediate level courses in machine learning and/or statistics.

Advances in Information Retrieval - Joemon M. Jose 2020-04-11

This two-volume set LNCS 12035 and 12036 constitutes the refereed proceedings of the 42nd European Conference on IR Research, ECIR 2020, held in Lisbon, Portugal, in April 2020.* The 55 full papers presented together with 8 reproducibility papers, 46 short papers, 10 demonstration papers, 12 invited CLEF papers, 7 doctoral consortium papers, 4 workshop papers, and 3 tutorials were carefully reviewed and selected from 457 submissions. They were organized in topical sections named: Part I: deep learning I; entities; evaluation; recommendation; information extraction; deep learning II; retrieval; multimedia; deep learning III; queries; IR - general; question answering, prediction, and bias; and deep learning IV. Part II: reproducibility papers; short papers; demonstration papers; CLEF organizers lab track; doctoral consortium papers; workshops; and tutorials. *Due to the COVID-19 pandemic, this conference was held virtually.

A Survey of Query Auto Completion in Information Retrieval - Fei Cai 2016-09-13

In information retrieval, query auto completion (QAC), also known as

type-ahead and auto-complete suggestion, helps users to formulate their query when they have an intent in mind but not a clear way of expressing this in a query. This monograph surveys this research.

Wsdm '18 - Wsdm 2018-10-25

We are delighted to welcome you to the Eleventh ACM International Conference on Web Search and Data Mining (WSDM 2018) held in Los Angeles, California, USA, on February 5-9, 2018. Now in its eleventh year, WSDM has become a top tier conference in Web-inspired research relating to search and data mining. As in previous years, we continued observing a growth in number of submissions. The conference this year, with 514 valid submissions, maintained the major boost that was observed last year in Cambridge UK (with 505 submissions as compared to 386 in year 2016 in San Francisco). Interestingly, we saw an increase in 10% in abstract submissions with a record breaking 757 abstracts (as compared to 687 in 2017) but only 68% of those ended up as valid submissions. This is partly due to a relatively high number of invalid submissions that did not adhere to our new double-blind review policy. The 514 valid submissions originated from 41 countries, out of which 84 (as compared to 80 in 2017) were accepted for full paper publication in the proceedings, thus reaching an acceptance rate of 16.12% (as compared to 15.94% last year) and within the range of the last 11 years (with a min of 15.5% and a max of 22.3%.) Unfortunately, three papers were withdrawn/rejected after acceptance due to their violations of WSDM guidelines. The final 81 papers that will be published in the proceedings are from 23 countries, spanning four continents, making this a truly international forum. Oral presentation slots were allocated to all papers. Yet, in order to maintain the single track model that most attendees prefer, we followed the spotlight short presentation plus poster approach that was introduced in 2012. Out of the 81 accepted papers, 58 were assigned such a two-minute spotlight slot, while 23 were assigned a long twenty-minute talk slot. The type of slot was chosen by the Senior PC members and Program co-chairs, mostly based on whether the topic and the content of the paper were best suited for a large group presentation or for a more focused and interactive poster style of presentation. The double-blind flavor we adopted this year allowed the authors to indicate the source of their data set, or deployment environment (so as not to refer to major commercial search engines as has often been done in the past). However, we observed that several authors failed to indicate their conflict of interest (COI) adequately and

we will make sure to enforce COI guidelines (through both awareness campaigns and hopefully automated tools) more rigorously now that the double-blind review policy has been adopted.

Test Collection Based Evaluation of Information Retrieval Systems - Mark Sanderson 2010-06-03

Use of test collections and evaluation measures to assess the effectiveness of information retrieval systems has its origins in work dating back to the early 1950s. Across the nearly 60 years since that work started, use of test collections is a de facto standard of evaluation. This monograph surveys the research conducted and explains the methods and measures devised for evaluation of retrieval systems, including a detailed look at the use of statistical significance testing in retrieval experimentation. This monograph reviews more recent examinations of the validity of the test collection approach and evaluation measures as well as outlining trends in current research exploiting query logs and live labs. At its core, the modern-day test collection is little different from the structures that the pioneering researchers in the 1950s and 1960s conceived of. This tutorial and review shows that despite its age, this long-standing evaluation method is still a highly valued tool for retrieval research.

Advances in Information Retrieval - Leif Azzopardi 2019-04-06

This two-volume set LNCS 11437 and 11438 constitutes the refereed proceedings of the 41st European Conference on IR Research, ECIR 2019, held in Cologne, Germany, in April 2019. The 48 full papers presented together with 2 keynote papers, 44 short papers, 8 demonstration papers, 8 invited CLEF papers, 11 doctoral consortium papers, 4 workshop papers, and 4 tutorials were carefully reviewed and selected from 365 submissions. They were organized in topical sections named: Modeling Relations; Classification and Search; Recommender Systems; Graphs; Query Analytics; Representation; Reproducibility (Systems); Reproducibility (Application); Neural IR; Cross Lingual IR; QA and Conversational Search; Topic Modeling; Metrics; Image IR; Short Papers; Demonstration Papers; CLEF Organizers Lab Track; Doctoral Consortium Papers; Workshops; and Tutorials.

An Introduction to Neural Information Retrieval - Bhaskar Mitra 2018-12-23

Efficient Query Processing for Scalable Web Search will be a valuable reference for researchers and developers working on This tutorial provides an accessible, yet comprehensive, overview of the state-of-the-art of Neural Information Retrieval.