Learning To Rank For Information Retrieval And Natural Language Processing Hang Li

The use of machine learning in mechanics is booming. Algorithms inspired by developments in the field of artificial intelligence today cover increasingly varied fields of application. This book illustrates recent results on coupling machine learning with computational mechanics, particularly for the construction of surrogate models or reduced order models. The articles contained in this compilation were presented at the EUROMECH Colloquium 597, « Reduced Order Modeling in Mechanics of Materials », held in Bad Herrenalb, Germany, from August 28th to August 31th 2018. In this book, Artificial Neural Networks are coupled to physics-based models. The tensor format of simulation data is exploited in surrogate models or for data pruning. Various reduced order models are proposed via machine learning strategies applied to simulation data. Since reduced order models have specific approximation errors, error estimators are also proposed in this book. The proposed numerical examples are very close to engineering problems. The reader would find this book to be a useful reference in identifying progress in machine learning and reduced order modeling for computational mechanics.

The controversial journalistic analysis of the mentality that fostered the Holocaust, from the author of The Origins of Totalitarianism Sparking a flurry of heated debate, Hannah Arendt’s authoritative and stunning report on the trial of German Nazi leader Adolf Eichmann first appeared as a series of articles in The New Yorker in 1963. This revised edition includes material that came to light after the trial, as well as Arendt’s postscript directly addressing the controversy that arose over her account. A major journalistic triumph by an intellectual of singular influence, Eichmann in Jerusalem is as shocking as it is informative—an unflinching look at one of the most unsettling (and unsettled) issues of the twentieth century.

This three-volume set LNAI 6911, LNAI 6912, and LNAI 6913 constitutes the refereed proceedings of the European conference on Machine Learning and Knowledge Discovery in Databases: ECML PKDD 2011, held in Athens, Greece, in September 2011. The 121 revised full papers presented together with 10 invited talks and 11 demos in the three volumes, were carefully reviewed and selected from about 600 paper submissions. The papers address all areas related to machine learning and knowledge discovery in databases as well as other innovative application domains such as supervised and unsupervised learning with some innovative contributions in fundamental issues; dimensionality reduction, distance and similarity learning, model learning and matrix/tensor analysis; graph mining, graphical models, hidden markov models, kernel methods, active and ensemble learning, semi-supervised and transductive learning, mining sparse representations, model learning, inductive logic programming, and statistical learning. a significant part of the papers covers novel and timely applications of data mining and machine learning in industrial domains.

A brilliant, authoritative, and fascinating history of America’s most puzzling era, the years 1920 to 1933, when the U.S. Constitution was amended to restrict one of America’s favorite pastimes: drinking alcoholic beverages. From its start, America has been awash in drink. The sailing vessel that brought John Winthrop to the shores of the New World in 1630 carried more beer than water. By the 1820s, liquor flowed so

plentifully it was cheaper than tea. That Americans would ever agree to relinquish their booze was as improbable as it was astonishing. Yet we did, and Last Call is Daniel Okrent's dazzling explanation of why we did it, what life under Prohibition was like, and how such an unprecedented degree of government interference in the private lives of Americans changed the country forever. Writing with both wit and historical acuity, Okrent reveals how Prohibition marked a confluence of diverse forces: the growing political power of the women's suffrage movement, which allied itself with the antiliquor campaign; the fear of small-town, native-stock Protestants that they were losing control of their country to the immigrants of the large cities; the anti-German sentiment stoked by World War I; and a variety of other unlikely factors, ranging from the rise of the automobile to the advent of the income tax. Through it all, Americans kept drinking, going to remarkably creative lengths to smuggle, sell, conceal, and convivially (and sometimes fatally) imbibe their favorite intoxicants. Last Call is peopled with vivid characters of an astonishing variety: Susan B. Anthony and Billy Sunday, William Jennings Bryan and bootlegger Sam Bronfman, Pierre S. du Pont and H. L. Mencken, Meyer Lansky and the incredible—if long-forgotten—federal official Mabel Walker Willebrandt, who throughout the twenties was the most powerful woman in the country. (Perhaps most surprising of all is Okrent's account of Joseph P. Kennedy's legendary, and long-misunderstood, role in the liquor business.) It's a book rich with stories from nearly all parts of the country. Okrent's narrative runs through smoky Manhattan speakeasies, where relations between the sexes were changed forever; California vineyards busily producing "sacramental" wine; New England fishing communities that gave up fishing for the more lucrative rum-running business; and in Washington, the halls of Congress itself, where politicians who had voted for Prohibition drank openly and without apology. Last Call is capacious, meticulous, and thrillingly told. It stands as the most complete history of Prohibition ever written and confirms Daniel Okrent's rank as a major American writer.

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

This book constitutes the refereed proceedings of the 33rd annual European Conference on Information Retrieval Research, ECIR 2011, held in Dublin, Ireland, in April 2010. The 45 revised full papers presented together with 24 poster papers, 17 short papers, and 6 tool demonstrations were carefully reviewed and selected from 223

full research paper submissions and 64 poster/demo submissions. The papers are organized in topical sections on text categorization, recommender systems, Web IR, IR evaluation, IR for Social Networks, cross-language IR, IR theory, multimedia IR, IR applications, interactive IR, and question answering /NLP.

A new edition of a graduate-level machine learning textbook that focuses on the analysis and theory of algorithms. This book is a general introduction to machine learning that can serve as a textbook for graduate students and a reference for researchers. It covers fundamental modern topics in machine learning while providing the theoretical basis and conceptual tools needed for the discussion and justification of algorithms. It also describes several key aspects of the application of these algorithms. The authors aim to present novel theoretical tools and concepts while giving concise proofs even for relatively advanced topics. Foundations of Machine Learning is unique in its focus on the analysis and theory of algorithms. The first four chapters lay the theoretical foundation for what follows; subsequent chapters are mostly self-contained. Topics covered include the Probably Approximately Correct (PAC) learning framework; generalization bounds based on Rademacher complexity and VC-dimension; Support Vector Machines (SVMs); kernel methods; boosting; on-line learning; multi-class classification; ranking; regression; algorithmic stability; dimensionality reduction; learning automata and languages; and reinforcement learning. Each chapter ends with a set of exercises. Appendixes provide additional material including concise probability review. This second edition offers three new chapters, on model selection, maximum entropy models, and conditional entropy models. New material in the appendixes includes a major section on Fenchel duality, expanded coverage of concentration inequalities, and an entirely new entry on information theory. More than half of the exercises are new to this edition.

Table of contents

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.

Deep Learning through Sparse Representation and Low-Rank Modeling bridges classical sparse and low rank models-those that emphasize problem-specific Interpretability-with recent deep network models that have enabled a larger learning capacity and better utilization of Big Data. It shows how the toolkit of deep learning is closely tied with the sparse/low rank methods and algorithms, providing a rich variety of theoretical and analytic tools to guide the design and interpretation of deep learning models. The development of the theory and models is supported by a wide variety of applications in computer vision, machine learning, signal processing, and data mining.

This book will be highly useful for researchers, graduate students and practitioners working in the fields of computer vision, machine learning, signal processing, optimization and statistics. Combines classical sparse and low-rank models and algorithms with the latest advances in deep learning networks Shows how the structure and algorithms of sparse and low-rank methods improves the performance and interpretability of Deep Learning models Provides tactics on how to build and apply customized deep learning models for various applications

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.

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

The goal of machine learning is to program computers to use example data or past experience to solve a given problem. Many successful applications of machine learning exist already, including systems that analyze past sales data to predict customer behavior, optimize robot behavior so that a task can be completed using minimum resources, and extract knowledge from bioinformatics data. Introduction to Machine Learning is a comprehensive textbook on the subject, covering a broad array of topics not usually included in introductory machine learning texts. Subjects include supervised learning; Bayesian decision theory; parametric, semi-parametric, and nonparametric methods; multivariate analysis; hidden Markov models; reinforcement learning; kernel machines; graphical models; Bayesian estimation; and statistical testing.Machine learning is rapidly becoming a skill that computer science students must master before graduation. The third edition of Introduction to Machine Learning reflects this shift, with added support for beginners, including selected solutions for exercises and additional example data sets (with code available online). Other substantial changes include discussions of outlier detection; ranking algorithms for perceptrons and support vector machines; matrix decomposition and spectral methods; distance estimation; new kernel algorithms; deep learning in multilayered perceptrons; and the nonparametric approach to Bayesian methods. All learning algorithms are explained so that students can easily move from the equations in the book to a computer program. The book can be used by both advanced undergraduates and graduate students. It will also be of interest to professionals who are concerned with the application of machine learning methods.

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

This book constitutes the proceedings of the 35th European Conference on IR Research, ECIR 2013, held in Moscow, Russia, in March 2013. The 55 full papers, 38 poster papers and 10 demonstrations presented in this volume were carefully reviewed and selected from 287 submissions. The papers are organized in the following topical sections: user aspects; multimedia and cross-media IR; data mining; IR theory and formal models; IR system architectures; classification; Web; event detection; temporal IR, and microblog search. Also included are 4 tutorial and 2 workshop presentations.

The ubiquitous challenge of learning and decision-making from rank data arises in situations where intelligent systems collect preference and behavior data from humans, learn from the data, and then use the data to help humans make efficient, effective, and timely decisions. Often, such data are represented by rankings. This book surveys some recent progress toward addressing the challenge from the considerations of statistics, computation, and socio-economics. We will cover classical statistical models for rank data, including random utility models, distance-based models, and mixture models. We will discuss and compare classical and state of-the-art algorithms, such as algorithms based on Minorize-Majorization (MM), Expectation-Maximization (EM), Generalized Method-of-Moments (GMM), rank breaking, and tensor decomposition. We will also introduce principled Bayesian preference elicitation frameworks for collecting rank data. Finally, we will examine socio-economic aspects of statistically desirable decision-making mechanisms, such as Bayesian estimators. This book can be useful in three ways: (1) for theoreticians in statistics and machine learning to better understand the considerations and caveats of learning from rank data, compared to learning from other types of data, especially cardinal data; (2) for practitioners to apply algorithms covered by the book for sampling, learning, and aggregation; and (3) as a textbook for graduate students or advanced undergraduate students to learn about the field. This book requires that the reader has basic knowledge in probability, statistics, and algorithms. Knowledge in social choice would also help but is not required.

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.

Search is one of the most critical functionalities of an e-commerce site. Almost every e-commerce site provides a search box. A customer expresses her intent about a product or a category of products in the form of one or more keywords and enters those keywords in the search box. The underlying search engine for an e-commerce site takes those keywords as input and formulate a search query. It subsequently queries its index and retrieves a ranked list of relevant items. The customer then goes through these from the top one after another. She then can click one of those items and lands on its product view page. That page contains detail information about the product. The customer then can add that product to her shopping cart and can purchase it. The internals of an e-commerce search engine is mostly similar to a web search engine. However, the ranking problem for an e-commerce search engine has several unique challenges. In this thesis we unfold the complexities of designing, optimizing and evaluating an e-commerce search engine. We begin with the aspect of the evaluation of the ranking algorithms for an e-commerce search and provide guidelines for conducting online randomized controlled experiments on a large e-commerce site. In this regard, we discuss managing biases, understanding of the metrics and the population, and the use of right statistical tests. Second, we define a formal framework for designing learning to rank (LTR) algorithms for e-commerce search optimizing the ranking for relevance, revenue, and discovery. We define a measure for discovery and describe the importance of that for an e-commercebusiness. We then propose a practical algorithm for integrating a discovery mechanism with trained learning to rank model. We also describe an approach for evaluating ranking algorithms involving exploration using offline simulation and counterfactuals. We furthermore, design variants of exploration algorithms using our formal framework. We show that a class of such explore algorithms can be designed to be monotonic sub-modular. It is thus possible to construct simple greedy algorithms with theoretical guarantees that maximize the exploration minimizing the loss of relevance and revenue. We conduct simulation studies following our proposed evaluation methodologies to show the

effectiveness of our algorithms. Third, we address the problem of incorporating diversity in e-commerce search. We design a knapsack based semi-bandit optimization algorithm for simultaneously learning to diversify and maximizing the revenue. We show that the regret of the algorithm is similar to an existing algorithm for learning to diversify web search although our algorithm is much more straightforward and efficient to realize. We further show that improving diversity in e-commerce search increases the median customer lifetime value (CLV) for an e-commerce business. Fourth, we address the problem of multi-objective learning to rank. We use the LambdaMart algorithm to realize our multi-objective algorithms. LambdaMart algorithm is widely used in Industry, won some recent "learning to rank" challenges. The authors of the algorithm showed that it could empirically optimize any non-smooth information retrieval ranking metrics such as normalized discounted cumulative gradient (NDCG). It is thus an efficient choice for designing multi-objective learning to rank algorithm. We also extend the physics-based intuition behind the design of lambda gradients to construct a dynamic update for one of the algorithms. We, furthermore, design a second more general algorithm that maintains a fixed number of solutions at each iteration and use a utility function given by a decision maker for that. The algorithm selects the best path that maximizes the NDCG scores of multiple objectives based on the preferences used in the utility function once it completes all the iterations to generate the model. This algorithm can achieve Pareto optimality with the given utility function if such a solution exists. Fifth, we address the problem of quantification and visualization of the excess supply and unmet demand using the contents of thequeries and items. We show the impact of such content gap in search experience. We quantify the content gap defining a distance between the topic distribution of text of the items and the queries. E-commerce companies can use the insight obtained from such topical content gaps in formulating strategies for demand generation and assortment planning. Finally, we consider the problem of optimization of ranking for the recommender systems for two-sided marketplaces by constructing a two-layer based learning architecture. We show that such two-layer models work better compared to using one model. Although this thesis focuses on the problems in e-commerce search, the algorithms that we developed are general and can have applications in many other domains. We expect to see this thesis serves as a guideline for several new directions of research in information retrieval and machine learning. Notably, we believe this thesis can propel the research for the problems that require integrating online learning with learning to rank, multi-objective learning to rank and regression algorithms, variants of application-specific multi-armed bandit algorithms, and search and recommendation applications for any marketplaces. rank, expert search and opinion detection.

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.

This book constitutes the refereed proceedings of the 39th European Conference on IR Research, ECIR 2017, held in Aberdeen, UK, in April 2017. The 36 full papers and 47 poster papers presented together with 5 Abstracts, were carefully reviewed and selected from 248 submissions. Being the premier European forum for the presentation of new research results in the field of Information Retrieval, ECIR features a wide range of topics such as: IR Theory and Practice; Deep Learning and IR; Web and Social Media IR; User Aspects; IR System Architectures; Content Representation and Processing; Evaluation; Multimedia and Cross-Media IR; Applications.

One of Mark Cuban’s top reads for better understanding A.I. (inc.com, 2021) Your comprehensive entry-level guide to machine learning While machine learning expertise doesn’t quite mean you can create your own Turing Test-proof android—as in the movie Ex Machina—it is a form of artificial intelligence and one of the most exciting technological means of identifying opportunities and solving problems fast and on a large scale. Anyone who masters the principles of machine learning is mastering a big part of our tech future and opening up incredible new directions in careers that include fraud detection, optimizing search results, serving real-time ads, credit-scoring, building accurate and sophisticated pricing models—and way, way more. Unlike most machine learning books, the fully updated 2nd Edition of Machine Learning For Dummies doesn't assume you have years of experience using programming languages such as Python (R source is also included in a downloadable form with comments and explanations), but lets you in on the ground floor, covering the entry-level materials that will get you up and running building models you need to perform practical tasks. It takes a look at the underlying—and fascinating—math principles that power machine learning but also shows that you don't need to be a math whiz to build fun new tools and apply them to your work and study. Understand the history of AI and machine learning Work with Python 3.8 and TensorFlow 2.x (and R as a download) Build and test your own models Use the latest datasets, rather than the worn out data found in other books Apply machine learning to real problems Whether you want to learn for college or to enhance your business or career performance, this friendly beginner's guide is your best introduction to machine learning, allowing you to become quickly confident using this amazing and fast-developing technology that's impacting lives for the better all over the world.

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 indomains 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.

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.

This book includes the proceedings of the first workshop on Recommender Systems in Fashion 2019. It presents a state of the art view of the advancements within the field of recommendation systems with focused application to e-commerce, retail and fashion. The volume covers contributions from academic as well as industrial researchers active within this emerging new field. Recommender Systems are often used to solve different complex problems in this scenario, such as social fashion-based recommendations (outfits inspired by influencers), product recommendations, or size and fit recommendations. The impact of social networks and the influence that fashion influencers have on the choices people make for

shopping is undeniable. For instance, many people use Instagram to learn about fashion trends from top influencers, which helps them to buy similar or even exact outfits from the tagged brands in the post. When traced, customers' social behavior can be a very useful guide for online shopping websites, providing insights on the styles the customers are really interested in, and hence aiding the online shops in offering better recommendations and facilitating customers quest for outfits. Another well known difficulty with recommendation of similar items is the large quantities of clothing items which can be considered similar, but belong to different brands. Relying only on implicit customer behavioral data will not be sufficient in the coming future to distinguish between for recommendation that will lead to an item being purchased and kept, vs. a recommendation that might result in either the customer not following it, or eventually return the item. Finding the right size and fit for clothes is one of the major factors not only impacting customers purchase decision, but also their satisfaction from e-commerce fashion platforms. Moreover, fashion articles have important sizing variations. Finally, customer preferences towards perceived article size and fit for their body remain highly personal and subjective which influences the definition of the right size for each customer. The combination of the above factors leaves the customers alone to face a highly challenging problem of determining the right size and fit during their purchase journey, which in turn has resulted in having more than one third of apparel returns to be caused by not ordering the right article size. This challenge presents a huge opportunity for research in intelligent size and fit recommendation systems and machine learning solutions with direct impact on both customer satisfaction and business profitability.

Learning to Rank for Information RetrievalSpringer Science & Business Media

"The amount of digital data we produce every day far surpasses our ability to process this data, and finding useful information in this constant flow of data has become one of the major challenges of the 21st century. Search engines are one way of accessing large data collections. Their algorithms have evolved far beyond simply matching search queries to sets of documents. Today's most sophisticated search engines combine hundreds of relevance signals to provide the best possible results for each searcher. Current approaches for tuning the parameters of search engines can be highly effective. However, they typically require considerable expertise and manual effort. They rely on supervised learning to rank, meaning that they learn from manually annotated examples of relevant documents for given queries. Obtaining large quantities of sufficiently accurate manual annotations is becoming increasingly difficult, especially for personalized search, access to sensitive data, or search in settings that change over time. In this thesis, I develop new online learning to rank techniques, based on insights from reinforcement learning. In contrast to supervised approaches, these methods allow search engines to learn directly from users' interactions. User interactions can typically be observed easily and cheaply, and reflect the preferences of real users. Interpreting user interactions and learning from them is challenging, because they can be biased and noisy. The contributions of this thesis include a novel interleaved comparison method, called probabilistic interleave, that allows unbiased comparisons of search engine result rankings, and methods for learning quickly and effectively from the resulting relative feedback. The obtained analytical and experimental results show how search engines can effectively learn from user interactions."--Omslag.

Summary Relevant Search demystifies relevance work. Using Elasticsearch, it teaches you how to return engaging search results to your users, helping you understand and leverage the internals of Lucene-based search engines. Purchase of the print book includes a free eBook in PDF, Kindle, and ePub formats from Manning Publications. About the Technology Users are accustomed to and expect instant, relevant search results. To achieve this, you must master the search engine. Yet for many developers, relevance ranking is mysterious or confusing. About the Book Relevant Search demystifies the subject and shows you that a search engine

is a programmable relevance framework. You'll learn how to apply Elasticsearch or Solr to your business's unique ranking problems. The book demonstrates how to program relevance and how to incorporate secondary data sources, taxonomies, text analytics, and personalization. In practice, a relevance framework requires softer skills as well, such as collaborating with stakeholders to discover the right relevance requirements for your business. By the end, you'll be able to achieve a virtuous cycle of provable, measurable relevance improvements over a search product's lifetime. What's Inside Techniques for debugging relevance? Applying search engine features to real problems? Using the user interface to guide searchers? A systematic approach to relevance? A business culture focused on improving search About the Reader For developers trying to build smarter search with Elasticsearch or Solr. About the Authors Doug Turnbull is lead relevance consultant at OpenSource Connections, where he frequently speaks and blogs. John Berryman is a data engineer at Eventbrite, where he specializes in recommendations and search. Foreword author, Trey Grainger, is a director of engineering at CareerBuilder and author of Solr in Action. Table of Contents The search relevance problem Search under the hood Debugging your first relevance problem Taming tokens Basic multifield search Term-centric search Shaping the relevance function Providing relevance feedback Designing a relevance-focused search application The relevance-centered enterprise Semantic and personalized search

This compact book explores standard tools for text classification, and teaches the reader how to use machine learning to decide whether a e-mail is spam or ham (binary classification), based on raw data from The SpamAssassin Public Corpus. Of course, sometimes the items in one class are not created equally, or we want to distinguish among them in some meaningful way. The second part of the book will look at how to not only filter spam from our email, but also placing "more important" messages at the top of the queue. This is a curated excerpt from the upcoming book "Machine Learning for Hackers."

The main challenge in learning-to-rank for information retrieval is the difficulty to di- rectly optimize ranking measures to automatically construct a ranking model from training data. It is mainly due to the fact that the ranking measures are determined by the order of ranked documents rather than the specific values of ranking model scores, thus they are non-convex, nondifferentiable and discontinuous. To address this issue, listwise approaches have been proposed where loss functions are defined either by exploiting a probabilistic model or by optimizing upper bounds or smoothed approximations of ranking measures. Even though very promising results have been achieved, there is still a mismatch between target cost and optimization cost. In this work, we present a novel learning algorithm that directly optimizes the ranking measures without resorting to any upper bounds or approx- imations. Our approach is essentially an iterative greedy coordinate descent method in optimization. For each iteration, we only update one parameter along one coordinate with all others fixed. Since the ranking measure is a stepwise function of a single parameter, we exploit an exhaustive line search algorithm to locate the interval with the smallest ranking measure along each coordinate. We pick the coordinate that leads to the largest reduction of ranking measure. In order to determine the optimal value of the parameter for the selected coordinate, we construct a probabilistic framework for the permutation, and maximize the likelihood of top-m ranked documents. This iterative procedure is continued until conver- gence. We conduct experiments of five datasets selected from Microsoft LETOR datasets, our experimental results show that the proposed direct rank algorithm outperforms several well-known state-of-the-art ranking algorithms.

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.

A comprehensive and rigorous introduction for graduate students and researchers, with applications in sequential decision-making problems.

ThePaci?c-AsiaConferenceonKnowledgeDiscoveryandData Mining hasbeen 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 inter- tional conference in the areas of data mining (DM) and knowledge discovery in database (KDD). It provides an international forum for researchers and ind- try practitioners to share their new ideas, original research results and prac- cal development experiences from all KDD-related areas including data mining, data warehousing, machine learning, databases, statistics, knowledge acqui- tion and automatic scienti?c discovery, data visualization, causal induction and knowledge-based systems. ForPAKDD2009,wereceived338researchpapersfromvariouscountriesand regions in Asia, Australia, North America, South America, Europe, and Africa. Every submission was rigorously reviewed by at least three reviewers with a doubleblindprotocol.Theinitialresultswerediscussedamongthereviewersand ?nally judged by the ProgramCommittee Chairs. When there was a con?ict, an additionalreviewwasprovidedbytheProgramCommitteeChairs.TheProgram 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.

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.

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.

If you’re an experienced programmer interested in crunching data, this book will get you started with machine learning—a toolkit of algorithms that enables computers to train themselves to automate useful tasks. Authors Drew Conway and John Myles White help you understand machine learning and statistics tools through a series of hands-on case studies, instead of a traditional math-heavy presentation. Each chapter focuses on a specific problem in machine learning, such as classification, prediction, optimization, and recommendation. Using the R programming language, you’ll learn how to analyze sample datasets and write simple machine learning algorithms. Machine Learning for Hackers is ideal for programmers from any background, including business, government, and academic research. Develop a naïve Bayesian classifier to determine if an email is spam, based only on its text Use linear regression to predict the number of page views for the top 1,000 websites Learn optimization techniques by attempting to break a simple letter cipher Compare and contrast U.S. Senators statistically, based on their voting records Build a “whom to follow” recommendation system from Twitter data

Endorsed by top AI authors, academics and industry leaders, The Hundred-Page Machine

Learning Book is the number one bestseller on Amazon and the most recommended book for starters and experienced professionals alike.

Machine learning approaches for learning ranking functions have been generating much interest from both the web information retrieval community and the machine learning community recently. It has the promise of improved relevancy of search engines and reduced demand for manual parameter tuning. We focus on developing a regression framework for learning to rank with complex loss functions. More specifically, this framework first applies functional iterative or boosting algorithm to compute updates for a given loss function and then fit the updates with a standard regression base learner. We explore supervised learning methodology from machine learning, and we distinguish two types of relevance judgments used as the training data: (1) absolute relevance judgments arising from explicit labeling of query-document pairs; and (2) relative relevance judgments extracted from user click-throughs of search results or converted from the absolute relevance judgments. Within the framework, we propose three novel ranking algorithms and illustrate their application to web search ranking. The first one is to calibrate the existing point-wise(univariant) regression loss to incorporate query difference in terms of introducing nuisance parameters in the statistical models, and we present an alternating optimization method to simultaneously learn the retrieval function and the nuisance parameters. It is an improvement over the existing approach within the category of learning to rank using point-wise regression loss. The second is an extension of gradient boosting methods for point-wise regression loss to complex(multi-variant) loss functions. It is based on optimization of quadratic upper bounds of the loss functions which allows us to present a rigorous convergence analysis of the algorithm. We illustrate an application of this approach in pair-wise preference learning to rank for Web search by combining both preference data and labeled data. The third one is a list-wise approach based on minimum effort optimization that takes into account the entire training data within a query at each iteration. We tackle this optimization problem using functional iterative methods where the update at each iteration is computed by solving an isotonic regression problem. This more global approach results in faster convergency and signficantly improved performance of the learned ranking functions over existing state-of-the-art methods. Experimental results are carried out using both data sets obtained from a commercial search engine and widely used IR benchmarking data, namely OHSUMED and TREC. Our results show significant improvements of our proposed methods over existing state-of-the-art methods.

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.

A complete overview of distant automatic speech recognition The performance of conventional Automatic Speech Recognition (ASR) systems degrades dramatically as soon as the

microphone is moved away from the mouth of the speaker. This is due to a broad variety of effects such as background noise, overlapping speech from other speakers, and reverberation. While traditional ASR systems underperform for speech captured with far-field sensors, there are a number of novel techniques within the recognition system as well as techniques developed in other areas of signal processing that can mitigate the deleterious effects of noise and reverberation, as well as separating speech from overlapping speakers. Distant Speech Recognitionpresents a contemporary and comprehensive description of both theoretic abstraction and practical issues inherent in the distant ASR problem. Key Features: Covers the entire topic of distant ASR and offers practical solutions to overcome the problems related to it Provides documentation and sample scripts to enable readers to construct state-of-the-art distant speech recognition systems Gives relevant background information in acoustics and filter techniques, Explains the extraction and enhancement of classification relevant speech features Describes maximum likelihood as well as discriminative parameter estimation, and maximum likelihood normalization techniques Discusses the use of multi-microphone configurations for speaker tracking and channel combination Presents several applications of the methods and technologies described in this book Accompanying website with open source software and tools to construct state-of-the-art distant speech recognition systems This reference will be an invaluable resource for researchers, developers, engineers and other professionals, as well as advanced students in speech technology, signal processing, acoustics, statistics and artificial intelligence fields.