

A Survey of Query Auto Completion in Information Retrieval

Fei Cai

Science and Technology on Information
Systems Engineering Laboratory,
National University of Defense Technology,
Changsha, China
University of Amsterdam, The Netherlands
f.cai@uva.nl

Maarten de Rijke

University of Amsterdam, The Netherlands
derijke@uva.nl

now

the essence of knowledge

Boston — Delft

Foundations and Trends[®] in Signal Processing

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
United States
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is

F. Cai and M. de Rijke. *A Survey of Query Auto Completion in Information Retrieval*. Foundations and Trends[®] in Signal Processing, vol. 10, no. 4, pp. 273–363, 2016.

This Foundations and Trends[®] issue was typeset in L^AT_EX using a class file designed by Neal Parikh. Printed on acid-free paper.

ISBN: 978-1-68083-201-3

© 2016 F. Cai and M. de Rijke

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The ‘services’ for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

Foundations and Trends[®] in Signal Processing

Volume 10, Issue 4, 2016

Editorial Board

Editor-in-Chief

Yonina Eldar

Technion - Israel Institute of Technology
Israel

Editors

Robert M. Gray

Founding Editor-in-Chief
Stanford University

Pao-Chi Chang
NCU, Taiwan

Pamela Cosman
UC San Diego

Michelle Effros
Caltech

Yariv Ephraim
GMU

Alfonso Farina
Selex ES

Sadaoki Furui
Tokyo Tech

Georgios Giannakis
University of Minnesota

Vivek Goyal
Boston University

Sinan Gunturk
Courant Institute

Christine Guillemot
INRIA

Robert W. Heath, Jr.
UT Austin

Sheila Hemami

Northeastern University

Lina Karam
Arizona State U

Nick Kingsbury
University of Cambridge

Alex Kot
NTU, Singapore

Jelena Kovacevic
CMU

Geert Leus
TU Delft

Jia Li
Penn State

Henrique Malvar
Microsoft Research

B.S. Manjunath
UC Santa Barbara

Urbashi Mitra
USC

Björn Ottersten
KTH Stockholm

Vincent Poor
Princeton University

Anna Scaglione
UC Davis

Mihaela van der Shaar
UCLA

Nicholas D. Sidiropoulos
TU Crete

Michael Unser
EPFL

P. P. Vaidyanathan
Caltech

Ami Wiesel
Hebrew U

Min Wu
University of Maryland

Josiane Zerubia
INRIA

Editorial Scope

Topics

Foundations and Trends[®] in Signal Processing publishes survey and tutorial articles in the following topics:

- Adaptive signal processing
- Audio signal processing
- Biological and biomedical signal processing
- Complexity in signal processing
- Digital signal processing
- Distributed and network signal processing
- Image and video processing
- Linear and nonlinear filtering
- Multidimensional signal processing
- Multimodal signal processing
- Multirate signal processing
- Multiresolution signal processing
- Nonlinear signal processing
- Randomized algorithms in signal processing
- Sensor and multiple source signal processing, source separation
- Signal decompositions, subband and transform methods, sparse representations
- Signal processing for communications
- Signal processing for security and forensic analysis, biometric signal processing
- Signal quantization, sampling, analog-to-digital conversion, coding and compression
- Signal reconstruction, digital-to-analog conversion, enhancement, decoding and inverse problems
- Speech/audio/image/video compression
- Speech and spoken language processing
- Statistical/machine learning
- Statistical signal processing

Information for Librarians

Foundations and Trends[®] in Signal Processing, 2016, Volume 10, 4 issues. ISSN paper version 1932-8346. ISSN online version 1932-8354. Also available as a combined paper and online subscription.

A Survey of Query Auto Completion in Information Retrieval

Fei Cai

Science and Technology on Information
Systems Engineering Laboratory,
National University of Defense Technology,
Changsha, China
University of Amsterdam, The Netherlands
f.cai@uva.nl

Maarten de Rijke
University of Amsterdam, The Netherlands
derijke@uva.nl

Contents

1	Introduction	2
1.1	Motivation	2
1.2	Historical notes	4
1.3	Aims of this survey	6
1.4	Structure of this survey	6
2	Query Auto Completion	8
2.1	Problem formulation	8
2.2	Aims of query auto completion	10
2.3	A taxonomy of solutions to query auto completion	11
3	Heuristic Approaches to Query Auto Completion	14
3.1	Time-sensitive query auto completion models	14
3.2	User-centered query auto completion models	21
3.3	Summary	25
4	Learning-based Approaches to Query Auto Completion	26
4.1	Learning from time-related characteristics	27
4.2	Learning from user interactions	30
4.3	Summary	34
5	Evaluation	35

5.1	Datasets	35
5.2	Evaluation metrics	38
5.3	Main experimental findings	40
5.4	Summary	42
6	Efficiency and Robustness	43
6.1	Computational efficiency in query auto completion	43
6.2	Error-tolerance in query auto completion	47
6.3	Summary	48
7	Presentation and Interaction	49
7.1	Presenting query auto completions results	50
7.2	Interactions with query auto completion results	52
7.3	Summary	55
8	Related Tasks	56
8.1	Query suggestion	56
8.2	Query expansion	59
8.3	Query correction	61
9	Conclusions	65
9.1	Summary of the survey	65
9.2	Open research directions	66
9.3	Concluding remarks	70
	Glossary	71
	Acknowledgements	72
	References	73

Abstract

In information retrieval, query auto completion (QAC), also known as type-ahead [[Xiao et al., 2013](#), [Cai et al., 2014b](#)] and auto-complete suggestion [[Jain and Mishne, 2010](#)], refers to the following functionality: given a prefix consisting of a number of characters entered into a search box, the user interface proposes alternative ways of extending the prefix to a full query. Ranking query completions is a challenging task due to the limited length of prefixes entered by users, the large volume of possible query completions matching a prefix, and the broad range of possible search intents. In recent years, a large number of query auto completion approaches have been proposed that produce ranked lists of alternative query completions by mining query logs.


In this survey, we review work on query auto completion that has been published before 2016. We focus mainly on web search and provide a formal definition of the query auto completion problem. We describe two dominant families of approaches to the query auto completion problem, one based on heuristic models and the other based on learning to rank. We also identify dominant trends in published work on query auto completion, viz. the use of time-sensitive signals and the use of user-specific signals. We describe the datasets and metrics that are used to evaluate algorithms for query auto completion. We also devote a chapter to efficiency and a chapter to presentation and interaction aspects of query auto completion. We end by discussing related tasks as well as potential research directions to further the area.

1


Introduction

1.1 Motivation

Word completion is a user interface feature that offers the user a list of words after one or more letters have been typed. It has been around as a feature of word editors and command shells for nearly half a century. In information retrieval, word completion is best known in the form of query auto completion, which provides users with suggested queries as they begin to enter their query in the search box. The user's incomplete input is often called a *query prefix* and the suggested queries are often called *query completions*. As shown in Figure 1.1, these query completions often start with the characters that the user has entered into the search box (see Figure 1.1a) and are adjusted as the user continues typing (see Figure 1.1b). Query auto completion (QAC) helps users to formulate their query when they have an intent in mind but not a clear way of expressing in a query. It helps to avoid possible spelling mistakes, especially on devices with small screens. It has been reported that, when submitting English queries, using QAC by selecting suggested query completions has saved more than 50% keystrokes of global Yahoo! searchers in 2014 [Zhang et al., 2015]. In addition, cutting down the search duration of users implies a lower load on the search engine, which results in savings in machine resources and maintenance [Bar-Yossef and Kraus, 2011].



(a) A list of four query completions for the prefix “information.”



(b) An updated list of four query completions for the prefix “information re.”

Figure 1.1: Examples of query auto completion.

Importantly, the use of QAC is not limited to web search engines such as Baidu, Bing, Google, or Yandex. QAC has also become part of other common services for search tasks. Facebook¹ provides QAC for finding friends, typically providing four suggestions; Twitter² uses QAC for searching tweets, also offering up to four suggestions; QAC is employed for product search in online shopping stores such as Amazon,³ which often produce up to ten suggestions; and it assists users to formulate requests for specific videos in online video-sharing websites such as Youtube,⁴ where users also get up to ten suggestions. As with web search, users of QAC functionality in other settings are helped to find an ideal query to address a particular search task. Research indicates that using the suggested queries can greatly improve user satisfaction, especially for informational queries [Song et al., 2011b].

¹<https://www.facebook.com>

²<https://twitter.com>

³<http://www.amazon.com>

⁴<https://www.youtube.com>

1.2 Historical notes

Query auto completion is a particular form of word prediction: when a writer writes the first letter or letters of a word, a word predictor lists possible words as choices and if the intended word is listed amongst the choices, it can be selected. In a different form of word prediction, the most likely next words (instead of completions) are listed. The core idea goes back at least half a century: [Longuet-Higgins and Ortony \[1968\]](#) describe a method to help decrease the number of keystrokes needed in order to complete a word, in their case commands and identifiers entered by developers. Another important motivation was to help individuals with physical disabilities increase their typing speed. Early work on word prediction focused both on assessing the cognitive benefits of word prediction and on algorithmic developments. For instance, early studies found that reduction of keystrokes through word prediction is often offset by having to scan the list of predicted words and selecting the desired word [[Vanderheiden and Kelso, 1987](#)]. Displaying five suggested words in a vertical manner was found to provide a reasonable balance between keystrokes saving and cognitive load [[Swiffin et al., 1987](#)].

Early algorithmic work on word prediction, that is, work from the 1970s and 1980s, can be grouped into three classes: character predictors, word completers, and combinations of the two [[Darragh, 1988](#)]. Character predictors accelerate input by making likely letters faster and easier to select. Word completers support text input by offering words based on an initial prefix of one or more letters entered by the user. Combined systems do both. The *Reactive Keyboard* introduced by [Darragh et al. \[1990\]](#) is a combined approach that also incorporates many facilities of today's query auto completion methods, including the use of popular words and of task appropriate words, both as identified from past query behavior.

[Witten and Darragh \[1992\]](#) provide an extensive treatment of work on word completion as an assistive technology up to the early 1990s. The *Reactive Keyboard* mentioned above is a prototypical example for the period as it incorporates the functionality of word completion and accelerates typewritten communication with a computer. Its word completion is mostly based on a standard dictionary; later approaches extended it to use adaptive modeling, which produces completions based on previously entered text. Similarly, a user's previously entered text, i.e., the search context, is taken into account

for personalized query auto completion [Bar-Yossef and Kraus, 2011]. Furthermore, a user's long-term search history has been explored for query auto completion [Cai et al., 2016a], but it has not been incorporated into a word completion system. Basically, word completion relies on a special tree structure to match the input and its completions [Witten and Darragh, 1992, Darragh et al., 1990], which is similar to the data structure used in query auto completion. In addition, both word completion and query auto completion often build on simple lexical models based on n-grams, which can be directly extracted from a text corpus or query log [Cai and de Rijke, 2016b].

Sentence completion is a slightly different task: given an initial fragment of a sentence, identify the remaining part of the sentence that the user intends to write. This task setting is motivated by processes such as answering questions sent by email in a call center, or writing letters in an administrative environment [Grabski and Scheffer, 2004]. Grabski and Scheffer [2004], Bickel et al. [2005], Nandi and Jagadish [2007] study this task based on lexical statistics of text collections.

The use of word prediction for search, goes back to the Emacs text editor [Ciccarelli, 1978], which offered a single line of feedback, not a list of suggestions. In the early 2000s, Raskin [2000] advocated the use of incremental search in which a user gets instant feedback as they enter a query based on what may exist in the content search over what he called *delimited search*, which basically is a traditional search interface allowing search in three steps: a user submits a query, the system computes a result page, the user receives the result page. Since then, query auto completion has been used by internet browsers, development environments, web sites, desktop search, operating systems, databases, email clients, and search engines, which are the focus of this survey. Around the middle of the 2000s, QAC was being offered by a growing number of search engines. For instance, Google Suggest, which launched in 2004,⁵ was an early showcase of query auto completion in the search engine setting. And around the same time Kayak,⁶ a travel metasearch engine, completed location related queries with travel related suggestions such as airport codes.

⁵<http://looksgoodworkswell.blogspot.nl/2005/12/distracting-or-narrowing-looking.html>

⁶<https://www.kayak.com>

1.3 Aims of this survey

Because of the clear benefits of QAC, a considerable number of algorithmic approaches to QAC have been proposed in the past few years. Query logs have proved to be a key asset underlying most of the recent research. In this paper we survey this research. We focus on summarizing the literature on QAC and provide a general understanding of the wealth of QAC approaches that is currently available. Our contributions in this survey can be summarized as follows:

1. We provide researchers who are working on query auto completion or related problems in the field of information retrieval with a good overview and analysis of state-of-the-art QAC approaches. In particular, for new researchers to the field, the survey can serve as an introduction to the state-of-the-art.
2. We offer a comprehensive perspective on QAC approaches by presenting a taxonomy of existing solutions. In addition, we present solutions for QAC under different conditions such as available high-resolution query logs, in-depth user interactions with QAC using eye-tracking, and elaborate user engagements in a QAC process. We also discuss practical issues related to QAC.
3. We present a detailed discussion of core challenges and promising open directions in QAC.

Our focus in this survey is on effectiveness of QAC. We do, however, include chapters on efficiency (Chapter 6) and on presentation and interaction aspects of query auto completion (Chapter 7).

1.4 Structure of this survey

The rest of this survey is organized as follows. Chapter 2 describes the problem of QAC as well as some closely related tasks; Chapter 3 and Chapter 4 detail published QAC approaches based on heuristic models and on learning algorithms, respectively; Chapter 5 presents publicly available benchmarks as well as metrics for QAC evaluation. In Chapter 6 we discuss efficiency and robustness aspects of query auto completion. In Chapter 7 we include a survey of studies into interface and interaction aspects of query auto completion. We include a short overview of three areas of research that are closely

related to QAC, i.e., query suggestion, query expansion and query correction in Chapter 8. We conclude the survey in Chapter 9 and point out follow-up research directions.

References

- Rakesh Agrawal, Sreenivas Gollapudi, Alan Halverson, and Samuel Jeong. Diversifying search results. In *Proceedings of the 2nd ACM International Conference on Web Search and Data Mining*, WSDM '09, pages 5–14, New York, NY, USA, 2009. ACM.
- Alia Amin, Michiel Hildebrand, Jacco van Ossenbruggen, Vanessa Evers, and Lynda Hardman. Organizing suggestions in autocompletion interfaces. In *Proceedings of the 31st European Conference on Information Retrieval*, ECIR '09, pages 521–529, Berlin, Heidelberg, 2009. Springer-Verlag.
- Aris Anagnostopoulos, Luca Becchetti, Carlos Castillo, and Aristides Gionis. An optimization framework for query recommendation. In *Proceedings of the Third ACM International Conference on Web Search and Data Mining*, WSDM '10, pages 161–170, New York, NY, USA, 2010. ACM.
- Kevin Bache, David Newman, and Padhraic Smyth. Text-based measures of document diversity. In *Proceedings of the 19th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '13, pages 23–31, New York, NY, USA, 2013. ACM.
- Ricardo Baeza-Yates, Carlos Hurtado, and Marcelo Mendoza. Query recommendation using query logs in search engines. In *Proceedings of the 2004 International Conference on Current Trends in Database Technology*, EDBT '04, pages 588–596, Berlin, Heidelberg, 2004. Springer-Verlag.
- Jing Bai, Jian-Yun Nie, Guihong Cao, and Hugues Bouchard. Using query contexts in information retrieval. In *Proceedings of the 30th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '07, pages 15–22, New York, NY, USA, 2007. ACM.

- Ziv Bar-Yossef and Naama Kraus. Context-sensitive query auto-completion. In *Proceedings of the 20th International World Wide Web Conference, WWW '11*, pages 107–116, New York, NY, USA, 2011. ACM.
- Holger Bast and Ingmar Weber. Type less, find more: Fast autocompletion search with a succinct index. In *Proceedings of the 29th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '06*, pages 364–371, New York, NY, USA, 2006. ACM.
- Holger Bast, Alexandru Chitea, Fabian Suchanek, and Ingmar Weber. Ester: Efficient search on text, entities, and relations. In *Proceedings of the 30th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '07*, pages 671–678, New York, NY, USA, 2007. ACM.
- Holger Bast, Christian W. Mortensen, and Ingmar Weber. Output-sensitive autocompletion search. *Information Retrieval*, 11(4):269–286, 2008.
- Steven M. Beitzel, Eric C. Jensen, Ophir Frieder, Abdur Chowdhury, and Greg Pass. Surrogate scoring for improved metasearch precision. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '05*, pages 583–584, New York, NY, USA, 2005. ACM.
- Paul N. Bennett, Ryen W. White, Wei Chu, Susan T. Dumais, Peter Bailey, Fedor Borisjuk, and Xiaoyuan Cui. Modeling the impact of short- and long-term behavior on search personalization. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '12*, pages 185–194, New York, NY, USA, 2012. ACM.
- Sumit Bhatia, Debapriyo Majumdar, and Prasenjit Mitra. Query suggestions in the absence of query logs. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '11*, pages 795–804, New York, NY, USA, 2011. ACM.
- Abhijit Bhole and Raghavendra Udupa. On correcting misspelled queries in email search. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence, AAAI '15*, pages 4266–4267, New York, NY, USA, 2015. AAAI Press.
- Steffen Bickel, Peter Haider, and Tobias Scheffer. Learning to complete sentences. In *Proceedings of the 16th European Conference on Machine Learning, ECML '05*, pages 497–504, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- Alexey Borisov, Ilya Markov, Maarten de Rijke, and Pavel Serdyukov. A neural click model for web search. In *Proceedings of the 25th International World Wide Web Conference, WWW '16*, pages 531–541, New York, NY, USA, 2016. ACM.

- Andrei Broder, Peter Ciccolo, Evgeniy Gabrilovich, Vanja Josifovski, Donald Metzler, Lance Riedel, and Jeffrey Yuan. Online expansion of rare queries for sponsored search. In *Proceedings of the 18th International World Wide Web Conference*, WWW '09, pages 511–520, New York, NY, USA, 2009. ACM.
- Christopher J.C. Burges, Krysta M. Svore, Paul N. Bennett, Andrzej Pastusiak, and Qiang Wu. Learning to rank using an ensemble of lambda-gradient models. *Journal of Machine Learning Research*, 14:25–35, 2011.
- Georg Buscher, Andreas Dengel, and Ludger van Elst. Query expansion using gaze-based feedback on the subdocument level. In *Proceedings of the 31st International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '08, pages 387–394, New York, NY, USA, 2008. ACM.
- Fei Cai and Maarten de Rijke. Selectively personalizing query auto completion. In *Proceedings of the 39th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '16, pages 993–996, New York, NY, USA, 2016a. ACM.
- Fei Cai and Maarten de Rijke. Learning from homologous queries and semantically related terms for query auto completion. *Information processing and Management*, 52(4):628–643, 2016b.
- Fei Cai, Shangsong Liang, and Maarten de Rijke. Personalized document re-ranking based on bayesian probabilistic matrix factorization. In *Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '14, pages 835–838, New York, NY, USA, 2014a. ACM.
- Fei Cai, Shangsong Liang, and Maarten de Rijke. Time-sensitive personalized query auto-completion. In *Proceedings of the 23rd ACM Conference on Information and Knowledge Management*, CIKM '14, pages 1599–1608, New York, NY, USA, 2014b. ACM.
- Fei Cai, Shangsong Liang, and Maarten de Rijke. Prefix-adaptive and time-sensitive personalized query auto completion. *IEEE Transactions on Knowledge and Data Engineering*, 28(9):2452–2466, September 2016a.
- Fei Cai, Ridho Reinanda, and Maarten de Rijke. Diversifying query auto-completion. *ACM Transactions on Information Systems*, 34(4):Article 25, 2016b.
- Fei Cai, Shuaiqiang Wang, and Maarten de Rijke. Behavior-based personalization in web search. *Journal of the Association for Information Science and Technology*, 2016c. To appear.
- Huanhuan Cao, Daxin Jiang, Jian Pei, Qi He, Zhen Liao, Enhong Chen, and Hang Li. Context-aware query suggestion by mining click-through and session data. In *Proceedings of the 14th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '08, pages 875–883, New York, NY, USA, 2008. ACM.

- Claudio Carpineto and Giovanni Romano. A survey of automatic query expansion in information retrieval. *ACM Computing Surveys*, 44(1):1–50, 2012.
- Olivier Chapelle and Ya Zhang. A dynamic bayesian network click model for web search ranking. In *Proceedings of the 18th International World Wide Web Conference*, WWW '09, pages 1–10, New York, NY, USA, 2009. ACM.
- Olivier Chapelle, Donald Metzler, Ya Zhang, and Pierre Grinspan. Expected reciprocal rank for graded relevance. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, CIKM '09, pages 621–630, New York, NY, USA, 2009. ACM.
- Surajit Chaudhuri and Raghav Kaushik. Extending autocompletion to tolerate errors. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data*, SIGMOD '09, pages 707–718, New York, NY, USA, 2009. ACM.
- Qing Chen, Mu Li, and Ming Zhou. Improving query spelling correction using web search results. In *Proceedings of the 2007 Conference on Empirical Methods in Natural Language Processing*, EMNLP '07, pages 181–189, Stroudsburg, PA, USA, 2007. Association for Computational Linguistics.
- Steve Chien and Nicole Immorlica. Semantic similarity between search engine queries using temporal correlation. In *Proceedings of the 14th International World Wide Web Conference*, WWW '05, pages 2–11, New York, NY, USA, 2005. ACM.
- Paul Alexandru Chirita, Claudiu S. Firan, and Wolfgang Nejdl. Personalized query expansion for the web. In *Proceedings of the 30th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '07, pages 7–14, New York, NY, USA, 2007. ACM.
- Aleksandr Chuklin, Ilya Markov, and Maarten de Rijke. *Click Models for Web Search*. Synthesis Lectures on Information Concepts, Retrieval, and Services. Morgan & Claypool Publishers, San Rafael, CA, USA, August 2015.
- Karen Church and Nuria Oliver. Understanding mobile web and mobile search use in today's dynamic mobile landscape. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services*, MobileHCI '11, pages 67–76, New York, NY, USA, 2011. ACM.
- Eugene Ciccarelli. An introduction to the Emacs editor. Technical Report AIM-447, MIT Artificial Intelligence Laboratory, Cambridge, Massachusetts, 1978.
- Charles L. A. Clarke, Eugene Agichtein, Susan Dumais, and Ryen W. White. The influence of caption features on clickthrough patterns in web search. In *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '07, pages 135–142, New York, NY, USA, 2007. ACM.

- Charles L.A. Clarke, Maheedhar Kolla, Gordon V. Cormack, Olga Vechtomova, Azin Ashkan, Stefan Büttcher, and Ian MacKinnon. Novelty and diversity in information retrieval evaluation. In *Proceedings of the 31st International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '08, pages 659–666, New York, NY, USA, 2008. ACM.
- Charles L.A. Clarke, Maheedhar Kolla, and Olga Vechtomova. An effectiveness measure for ambiguous and underspecified queries. In *Proceedings of the 2009 International Conference on the Theory of Information Retrieval*, ICTIR '09, pages 188–199, Berlin, Heidelberg, 2009. Springer-Verlag.
- Keven Collins-Thompson, Paul N. Bennett, Ryen W. White, Sebastian de la Chica, and David Sontag. Personalizing web search results by reading level. In *Proceedings of the 20th ACM Conference on Information and Knowledge Management*, CIKM '11, pages 403–412, New York, NY, USA, 2011. ACM.
- Keven Collins-Thompson, Paul Bennett, Charles L.A. Clarke, and Ellen M. Voorhees. TREC 2013 web track overview. In *Proceedings of the 21st Text REtrieval Conference*, TREC '13, pages 1–15, Berlin, Heidelberg, 2013. Springer-Verlag.
- Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. The MIT Press, 3rd edition, 2009.
- Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. An experimental comparison of click position-bias models. In *Proceedings of the 1st ACM International Conference on Web Search and Data Mining*, WSDM '08, pages 87–94, New York, NY, USA, 2008. ACM.
- Nick Craswell, Rosie Jones, Georges Dupret, and Evelyne Viegas, editors. *WSCD '09: Proceedings 2009 Workshop on Web Search Click Data*, Proceedings of the 2009 Workshop on Web Search Click Data, 2009. ACM.
- Silviu Cucerzan and Eric Brill. Spelling correction as an iterative process that exploits the collective knowledge of web users. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, EMNLP '04, pages 293–300, Stroudsburg, PA, USA, 2004. Association for Computational Linguistics.
- Hang Cui, Ji-Rong Wen, Jian-Yun Nie, and Wei-Ying Ma. Probabilistic query expansion using query logs. In *Proceedings of the 11th International World Wide Web Conference*, WWW '02, pages 325–332, New York, NY, USA, 2002. ACM.
- Van Dang and Bruce W. Croft. Term level search result diversification. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '13, pages 603–612, New York, NY, USA, 2013. ACM.

- Van Dang and W. Bruce Croft. Diversity by proportionality: An election-based approach to search result diversification. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '12, pages 65–74, New York, NY, USA, 2012. ACM.
- John J. Darragh. Adaptive predictive text generation and the reactive keyboard. Master's thesis, Computer Science Department, University of Calgary, 1988.
- John J. Darragh, Ian H. Witten, and Mark L. James. The reactive keyboard: A predictive typing aid. *Computer*, 23(11):41–49, 1990.
- Giovanni Di Santo, Richard McCreadie, Craig Macdonald, and Iadh Ounis. Comparing approaches for query autocompletion. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '15, pages 775–778, New York, NY, USA, 2015. ACM.
- M. Benjamin Dias, Dominique Locher, Ming Li, Wael El-Deredy, and Paulo J. G. Lisboa. The value of personalised recommender systems to e-business: A case study. In *Proceedings of the 2008 ACM Conference on Recommender Systems*, RecSys '08, pages 291–294, New York, NY, USA, 2008. ACM.
- Zhicheng Dou, Ruihua Song, and Ji-Rong Wen. A large-scale evaluation and analysis of personalized search strategies. In *Proceedings of the 16th International World Wide Web Conference*, WWW '07, pages 581–590, New York, NY, USA, 2007. ACM.
- Huizhong Duan and Bo-June (Paul) Hsu. Online spelling correction for query completion. In *Proceedings of the 20th International World Wide Web Conference*, WWW '11, pages 117–126, New York, NY, USA, 2011. ACM.
- Huizhong Duan, Yanen Li, ChengXiang Zhai, and Dan Roth. A discriminative model for query spelling correction with latent structural SVM. In *Proceedings of the 2012 Conference on Empirical Methods in Natural Language Processing*, EMNLP '12, pages 1511–1521, Stroudsburg, PA, USA, 2012. Association for Computational Linguistics.
- Georges E. Dupret and Benjamin Piwowarski. A user browsing model to predict search engine click data from past observations. In *Proceedings of the 31st International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '08, pages 331–338, New York, NY, USA, 2008. ACM.
- Henry Feild and James Allan. Task-aware query recommendation. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '13, pages 83–92, New York, NY, USA, 2013. ACM.

- Daniel M. Fleder and Kartik Hosanagar. Recommender systems and their impact on sales diversity. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, EC '07, pages 192–199, New York, NY, USA, 2007. ACM.
- Bruno M. Fonseca, Paulo Golgher, Bruno Póssas, Berthier Ribeiro-Neto, and Nivio Ziviani. Concept-based interactive query expansion. In *Proceedings of the 14th ACM Conference on Information and Knowledge Management*, CIKM '05, pages 696–703, New York, NY, USA, 2005. ACM.
- Guillem Francès, Xiao Bai, B. Barla Cambazoglu, and Ricardo Baeza-Yates. Improving the efficiency of multi-site web search engines. In *Proceedings of the 7th ACM International Conference on Web Search and Data Mining*, WSDM '14, pages 3–12, New York, NY, USA, 2014. ACM.
- George W. Furnas, Thomas K. Landauer, Louis M. Gomez, and Susan T. Dumais. The vocabulary problem in human-system communication. *Communications of the ACM*, 30(11):964–971, 1987.
- Jianfeng Gao, Xiaolong Li, Daniel Micol, Chris Quirk, and Xu Sun. A large scale ranker-based system for search query spelling correction. In *Proceedings of the 23rd International Conference on Computational Linguistics*, COLING '10, pages 358–366, Stroudsburg, PA, USA, 2010. Association for Computational Linguistics.
- Nadav Golbandi, Liran Katzir, Yehuda Koren, and Ronny Lempel. Expediting search trend detection via prediction of query counts. In *Proceedings of the 6th ACM International Conference on Web Search and Data Mining*, WSDM '13, pages 295–304, New York, NY, USA, 2013. ACM.
- Paul Goodwin. The Holt-Winters approach to exponential smoothing: 50 years old and going strong. *Foresight: The International Journal of Applied Forecasting*, 1(19):30–33, 2010.
- Korinna Grabski and Tobias Scheffer. Sentence completion. In *Proceedings of the 27th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '04, pages 433–439, New York, NY, USA, 2004. ACM.
- Laura A. Granka, Thorsten Joachims, and Geri Gay. Eye-tracking analysis of user behavior in www search. In *Proceedings of the 27th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '04, pages 478–479, New York, NY, USA, 2004. ACM.
- Fan Guo, Chao Liu, and Yi Min Wang. Efficient multiple-click models in web search. In *Proceedings of the 2nd ACM International Conference on Web Search and Data Mining*, WSDM '09, pages 124–131, New York, NY, USA, 2009a. ACM.

- Jiafeng Guo, Gu Xu, Xueqi Cheng, and Hang Li. Named entity recognition in query. In *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '09, pages 267–274, New York, NY, USA, 2009b. ACM.
- Jiafeng Guo, Xueqi Cheng, Gu Xu, and Huawei Shen. A structured approach to query recommendation with social annotation data. In *Proceedings of the 19th ACM Conference on Information and Knowledge Management*, CIKM '10, pages 619–628, New York, NY, USA, 2010. ACM.
- Jiafeng Guo, Xueqi Cheng, Gu Xu, and Xiaofei Zhu. Intent-aware query similarity. In *Proceedings of the 20th ACM Conference on Information and Knowledge Management*, CIKM '11, pages 259–268, New York, NY, USA, 2011. ACM.
- Ben He and Iadh Ounis. Finding good feedback documents. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, CIKM '09, pages 2011–2014, New York, NY, USA, 2009. ACM.
- Jiyin He, Martha Larson, and Maarten de Rijke. Using coherence-based measures to predict query difficulty. In *30th European Conference on Information Retrieval (ECIR 2008)*, number 4956 in LNCS, pages 689–694. Springer, 2008.
- Qi He, Daxin Jiang, Zhen Liao, Steven C. H. Hoi, Kuiyu Chang, Ee-Peng Lim, and Hang Li. Web query recommendation via sequential query prediction. In *Proceedings of the 25th International Conference on Data Engineering*, ICDE '09, pages 1443–1454, Washington, DC, USA, 2009. IEEE Computer Society.
- Marti A. Hearst. *Search User Interfaces*. Cambridge University Press, 2009.
- Katja Hofmann, Shimon Whiteson, and Maarten de Rijke. A probabilistic method for inferring preferences from clicks. In *Proceedings of the 20th ACM Conference on Information and Knowledge Management*, CIKM '11. ACM, October 2011.
- Katja Hofmann, Anne Schuth, Shimon Whiteson, and Maarten de Rijke. Reusing historical interaction data for faster online learning to rank for IR. In *Proceedings of the Sixth ACM International Conference on Web Search and Data Mining*, WSDM '13, pages 183–192, New York, NY, USA, 2013a. ACM.
- Katja Hofmann, Shimon Whiteson, and Maarten de Rijke. Balancing exploration and exploitation in listwise and pairwise online learning to rank for information retrieval. *Information Retrieval*, 16(1):63–90, 2013b.
- Katja Hofmann, Bhaskar Mitra, Filip Radlinski, and Milad Shokouhi. An eye-tracking study of user interactions with query auto completion. In *Proceedings of the 23rd ACM Conference on Information and Knowledge Management*, CIKM '14, pages 549–558, New York, NY, USA, 2014. ACM.
- Charles C. Holt. Forecasting seasonals and trends by exponentially weighted moving averages. *International Journal of Forecasting*, 20(1):5–10, 2004.

- Bo-June (Paul) Hsu and Giuseppe Ottaviano. Space-efficient data structures for top-k completion. In *Proceedings of the 22nd International World Wide Web Conference*, WWW '13, pages 583–594, New York, NY, USA, 2013. ACM.
- Bouke Huurnink, Laura Hollink, Wietske van den Heuvel, and Maarten de Rijke. Search behavior of media professionals at an audiovisual archive: A transaction log analysis. *Journal of the Association for Information Science and Technology*, 61(6):1180–1197, 2010.
- Alpa Jain and Gilad Mishne. Organizing query completions for web search. In *Proceedings of the 19th ACM Conference on Information and Knowledge Management*, CIKM '10, pages 1169–1178, New York, NY, USA, 2010. ACM.
- Kalervo Järvelin and Jaana Kekäläinen. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems*, 20(4):422–446, 2002.
- Shengyue Ji, Guoliang Li, Chen Li, and Jianhua Feng. Efficient interactive fuzzy keyword search. In *Proceedings of the 18th International World Wide Web Conference*, WWW '09, pages 371–380, New York, NY, USA, 2009. ACM.
- Di Jiang, Kenneth Wai-Ting Leung, and Wilfred Ng. Context-aware search personalization with concept preference. In *Proceedings of the 20th ACM Conference on Information and Knowledge Management*, CIKM '11, pages 563–572, New York, NY, USA, 2011. ACM.
- Jyun-Yu Jiang, Yen-Yu Ke, Pao-Yu Chien, and Pu-Jen Cheng. Learning user reformulation behavior for query auto-completion. In *Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '14, pages 445–454, New York, NY, USA, 2014. ACM.
- Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. Accurately interpreting clickthrough data as implicit feedback. In *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '05, pages 154–161, New York, NY, USA, 2005. ACM.
- Hideo Joho and Joemon M. Jose. A comparative study of the effectiveness of search result presentation on the web. In *Proceedings of the 28th European Conference on Information Retrieval*, ECIR '06, pages 302–313, Berlin, Heidelberg, 2006. Springer-Verlag.
- Hideo Joho and Joemon M. Jose. Effectiveness of additional representations for the search result presentation on the web. *Information Processing and Management*, 44(1):226–241, 2008.

- Rosie Jones and Kristina Lisa Klinkner. Beyond the session timeout: Automatic hierarchical segmentation of search topics in query logs. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management, CIKM '08*, pages 699–708, New York, NY, USA, 2008. ACM.
- Maryam Kamvar and Shumeet Baluja. A large scale study of wireless search patterns: Google mobile search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '06*, pages 701–709, New York, NY, USA, 2006. ACM.
- Maryam Kamvar and Shumeet Baluja. The role of context in query input: Using contextual signals to complete queries on mobile devices. In *Proceedings of the 11th International Conference on Human Computer Interaction with Mobile Devices and Services, MobileHCI '09*, pages 405–412, New York, NY, USA, 2009. ACM.
- Dimitrios Kastrinakis and Yannis Tzitzikas. Advancing search query auto-completion services with more and better suggestions. In *Proceedings of the 10th International Conference on Web Engineering, ICWE'10*, pages 35–49, Berlin, Heidelberg, 2010. Springer-Verlag.
- Makoto P. Kato, Tetsuya Sakai, and Katsumi Tanaka. Structured query suggestion for specialization and parallel movement: Effect on search behaviors. In *Proceedings of the 21st International Conference on World Wide Web, WWW '12*, pages 389–398, New York, NY, USA, 2012. ACM.
- Makoto P. Kato, Tetsuya Sakai, and Katsumi Tanaka. When do people use query suggestion? A query suggestion log analysis. *Information Retrieval*, 16(6):725–746, 2013.
- Gabriella Kazai, Jaap Kamps, Marijn Koolen, and Natasa Milic-Frayling. Crowdsourcing for book search evaluation: Impact of hit design on comparative system ranking. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '11*, pages 205–214, New York, NY, USA, 2011. ACM.
- Eugene Kharitonov, Craig Macdonald, Pavel Serdyukov, and Iadh Ounis. User model-based metrics for offline query suggestion evaluation. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '13*, pages 633–642, New York, NY, USA, 2013. ACM.
- Youngho Kim and W. Bruce Croft. Diversifying query suggestions based on query documents. In *Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '14*, pages 891–894, New York, NY, USA, 2014. ACM.

- Anagha Kulkarni, Jaime Teevan, Krysta M. Svore, and Susan T. Dumais. Understanding temporal query dynamics. In *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, WSDM '11*, pages 167–176, New York, NY, USA, 2011. ACM.
- Kyung Soon Lee, W. Bruce Croft, and James Allan. A cluster-based resampling method for pseudo-relevance feedback. In *Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '08*, pages 235–242, New York, NY, USA, 2008. ACM.
- Damien Lefortier, Pavel Serdyukov, and Maarten de Rijke. Online exploration for detecting shifts in fresh intent. In *Proceedings of the 23rd ACM Conference on Information and Knowledge Management, CIKM '14*. ACM, November 2014.
- Xin Lei, Andrew Senior, Alexander Gruenstein, and Jeffrey Sorensen. Accurate and compact large vocabulary speech recognition on mobile devices. In *Interspeech 2013*, pages 662–665, 2013.
- Guoliang Li, Shengyue Ji, Chen Li, and Jianhua Feng. Efficient fuzzy full-text type-ahead search. *The VLDB Journal*, 20(4):617–640, 2011.
- Liangda Li, Hongbo Deng, Anlei Dong, Yi Chang, Hongyuan Zha, and Ricardo Baeza-Yates. Analyzing user's sequential behavior in query auto-completion via Markov processes. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '15*, pages 123–132, New York, NY, USA, 2015. ACM.
- Xinyi Li, Bob Schijvenaars, and Maarten de Rijke. Investigating search behavior and search failures in academic search. In *Submitted*, 2016.
- Yanen Li, Huizhong Duan, and ChengXiang Zhai. A generalized hidden Markov model with discriminative training for query spelling correction. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '12*, pages 611–620, New York, NY, USA, 2012. ACM.
- Yanen Li, Anlei Dong, Hongning Wang, Hongbo Deng, Yi Chang, and ChengXiang Zhai. A two-dimensional click model for query auto-completion. In *Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '14*, pages 455–464, New York, NY, USA, 2014. ACM.
- Zhen Liao, Daxin Jiang, Enhong Chen, Jian Pei, Huanhuan Cao, and Hang Li. Mining concept sequences from large-scale search logs for context-aware query suggestion. *ACM Transactions on Intelligent Systems and Technology*, 3(1):Article 17, 2011.

- Thomas Lin, Patrick Pantel, Michael Gamon, Anitha Kannan, and Ariel Fuxman. Active objects: Actions for entity-centric search. In *Proceedings of the 21st International World Wide Web Conference, WWW '12*, pages 589–598, New York, NY, USA, 2012. ACM.
- Chao Liu, Fan Guo, and Christos Faloutsos. Bayesian browsing model: Exact inference of document relevance from petabyte-scale data. *ACM Transactions on Knowledge Discovery from Data*, 4(4):19:1–19:26, 2010a.
- Chao Liu, Ryen W. White, and Susan Dumais. Understanding web browsing behaviors through weibull analysis of dwell time. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '10*, pages 379–386, New York, NY, USA, 2010b. ACM.
- Ning Liu, Jun Yan, Shuicheng Yan, Weiguo Fan, and Zheng Chen. Web query prediction by unifying model. In *Proceedings of the 2008 IEEE International Conference on Data Mining Workshops, ICDMW '08*, pages 437–441, Washington, DC, USA, 2008. IEEE Computer Society.
- Shuang Liu, Fang Liu, Clement Yu, and Weiyi Meng. An effective approach to document retrieval via utilizing wordnet and recognizing phrases. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '04*, pages 266–272, New York, NY, USA, 2004. ACM.
- Tie-Yan Liu. Learning to rank for information retrieval. *Foundations and Trends in Information Retrieval*, 3(3):225–331, 2009.
- Yang Liu, Ruihua Song, Yu Chen, Jian-Yun Nie, and Ji-Rong Wen. Adaptive query suggestion for difficult queries. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '12*, pages 15–24, New York, NY, USA, 2012. ACM.
- Yiqun Liu, Ruihua Song, Min Zhang, Zhicheng Dou, Takehiro Yamamoto, Makoto Kato, Hiroaki Ohshima, and Ke Zhou. Overview of the NTCIR-11 IMine task. In *Proceedings of the 11th NTCIR Conference, NTCIR '11*, Tokyo, Japan, 2014.
- Ziyang Liu, Sivaramakrishnan Natarajan, and Yi Chen. Query expansion based on clustered results. *Proceedings of the VLDB Endowment*, 4(6):350–361, 2011.
- H. Christopher Longuet-Higgins and Andrew Ortony. The adaptive memorization of sequences. In *Proceedings of the 3rd Annual Machine Intelligence Workshop, Machine Intelligence 3*, pages 311–322. Edinburgh University Press, 1968.
- Claudio Lucchese, Salvatore Orlando, Raffaele Perego, Fabrizio Silvestri, and Gabriele Tolomei. Identifying task-based sessions in search engine query logs. In *Proceedings of the 4th ACM International Conference on Web Search and Data Mining, WSDM '11*, pages 277–286, New York, NY, USA, 2011. ACM.

- Hao Ma, Haixuan Yang, Irwin King, and Michael R. Lyu. Learning latent semantic relations from clickthrough data for query suggestion. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management, CIKM '08*, pages 709–718, New York, NY, USA, 2008. ACM.
- Hao Ma, Michael R. Lyu, and Irwin King. Diversifying query suggestion results. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence, AAAI '10*, pages 1399–1404, New York, NY, USA, 2010. AAAI Press.
- Zhongrui Ma, Yu Chen, Ruihua Song, Tetsuya Sakai, Jiaheng Lu, and Ji-Rong Wen. New assessment criteria for query suggestion. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '12*, pages 1109–1110, New York, NY, USA, 2012. ACM.
- Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze. *Introduction to Information Retrieval*. Cambridge University Press, New York, NY, USA, 2008.
- Gary Marchionini and Ryen White. Find what you need, understand what you find. *International Journal of Human-Computer Interaction*, 23(3):205–237, 2007.
- Dhruv Matani. An $O(k \log n)$ algorithm for prefix based ranked autocomplete. <http://www.dhruvbird.com/autocomplete.pdf>, 2011.
- Nicolaas Matthijs and Filip Radlinski. Personalizing web search using long term browsing history. In *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, WSDM '11*, pages 25–34, New York, NY, USA, 2011. ACM.
- Qiaozhu Mei, Dengyong Zhou, and Kenneth Church. Query suggestion using hitting time. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management, CIKM '08*, pages 469–478, New York, NY, USA, 2008. ACM.
- Edgar Meij, Peter Mika, and Hugo Zaragoza. An evaluation of entity and frequency based query completion methods. In *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '09*, pages 678–679, New York, NY, USA, 2009. ACM.
- Edgar Meij, Marc Bron, Laura Hollink, Bouke Huurnink, and Maarten de Rijke. Mapping queries to the linking open data cloud: A case study using DBpedia. *Journal of Web Semantics*, 9(4):418–433, 2011.
- Vlachos Michail, Meek Christopher, Vagena Zografoula, and Gunopulos Dimitrios. Identifying similarities periodicities and bursts for online search queries. In *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data, SIGMOD '04*, pages 131–142, New York, NY, USA, 2004. ACM.
- Gilad Mishne and Maarten de Rijke. A study of blog search. In *Proceedings 28th European Conference on Information Retrieval (ECIR 2006)*, number 3936 in LNCS, pages 289–301. Springer, April 2006.

- Bhaskar Mitra. Exploring session context using distributed representations of queries and reformulations. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '15, pages 3–12, New York, NY, USA, 2015. ACM.
- Bhaskar Mitra and Nick Craswell. Query auto-completion for rare prefixes. In *Proceedings of the 24th ACM Conference on Information and Knowledge Management*, CIKM '15, pages 1755–1758, New York, NY, USA, 2015. ACM.
- Bhaskar Mitra, Milad Shokouhi, Filip Radlinski, and Katja Hofmann. On user interactions with query auto-completion. In *Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '14, pages 1055–1058, New York, NY, USA, 2014. ACM.
- Meredith Ringel Morris, Jaime Teevan, and Steve Bush. Enhancing collaborative web search with personalization: Groupization, smart splitting, and group highlighting. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work*, CSCW '08, pages 481–484, New York, NY, USA, 2008. ACM.
- Arnab Nandi and H. V. Jagadish. Effective phrase prediction. In *Proceedings of the 33rd International Conference on Very Large Data Bases*, VLDB '07, pages 219–230, Berlin, Heidelberg, 2007. VLDB Endowment.
- Harrie Oosterhuis, Anne Schuth, and Maarten de Rijke. Probabilistic multileave gradient descent. In *ECIR 2016: 38th European Conference on Information Retrieval*, LNCS, pages 661–668. Springer, 2016.
- Greg Pass, Abdur Chowdhury, and Cayley Torgeson. A picture of search. In *Proceedings of the 1st International Conference on Scalable Information Systems*, InfoScale '06, pages 1–7, New York, NY, USA, 2006. ACM.
- Jef Raskin. *The Humane Interface*. Addison-Wesley Professional, 2000.
- Ridho Reinanda, Edgar Meij, and Maarten de Rijke. Mining, ranking and recommending entity aspects. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '15, pages 263–272, New York, NY, USA, 2015. ACM.
- Matthew Richardson, Ewa Dominowska, and Robert Ragno. Predicting clicks: Estimating the click-through rate for new ads. In *Proceedings of the 16th International World Wide Web Conference*, WWW '07, pages 521–530, New York, NY, USA, 2007. ACM.
- Soo Young Rieh and Hong Xie. Analysis of multiple query reformulations on the web: The interactive information retrieval context. *Information Processing and Management*, 42(3):751 – 768, 2006.

- Stefan Riezler, Yi Liu, and Alexander Vasserman. Translating queries into snippets for improved query expansion. In *Proceedings of the 22nd International Conference on Computational Linguistics, COLING '08*, pages 737–744, Stroudsburg, PA, USA, 2008. Association for Computational Linguistics.
- Patrick Ruch. Using contextual spelling correction to improve retrieval effectiveness in degraded text collections. In *Proceedings of the 19th International Conference on Computational Linguistics, COLING '02*, pages 1–7, Stroudsburg, PA, USA, 2002. Association for Computational Linguistics.
- Tetsuya Sakai, Zhicheng Dou, Takehiro Yamamoto, Yiqun Liu, Min Zhang, and Ruihua Song. Overview of the NTCIR-10 INTENT-2 task. In *Proceedings of the 10th NTCIR Conference, NTCIR '10*, pages 94–123, Tokyo, Japan, 2013.
- Rodrygo L. T. Santos, Craig Macdonald, and Iadh Ounis. Learning to rank query suggestions for adhoc and diversity search. *Information Retrieval*, 16:429–451, 2013.
- Anne Schuth, Harrie Oosterhuis, Shimon Whiteson, and Maarten de Rijke. Multileave gradient descent for fast online learning to rank. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining, WSDM '16*, pages 457–466, New York, NY, USA, 2016. ACM.
- Xuehua Shen, Bin Tan, and ChengXiang Zhai. Context-sensitive information retrieval using implicit feedback. In *Proceedings of the 28th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '05*, pages 43–50, New York, NY, USA, 2005. ACM.
- Yelong Shen, Xiaodong He, Jianfeng Gao, Li Deng, and Gregoire Mesnil. Learning semantic representations using convolutional neural networks for web search. In *Proceedings of the 23rd International World Wide Web Conference, WWW '14*, pages 373–374, New York, NY, USA, 2014. ACM.
- Milad Shokouhi. Detecting seasonal queries by time-series analysis. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '11*, pages 1171–1172, New York, NY, USA, 2011. ACM.
- Milad Shokouhi. Learning to personalize query auto-completion. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '13*, pages 103–112, New York, NY, USA, 2013. ACM.
- Milad Shokouhi and Kira Radinsky. Time-sensitive query auto-completion. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '12*, pages 601–610, New York, NY, USA, 2012. ACM.

- Milad Shokouhi, Marc Sloan, Paul N. Bennett, Kevyn Collins-Thompson, and Siranush Sarkizova. Query suggestion and data fusion in contextual disambiguation. In *Proceedings of the 24th International World Wide Web Conference, WWW '15*, pages 971–980, New York, NY, USA, 2015. ACM.
- Ruihua Song, Min Zhang, Tetsuya Sakai, Makoto P. Kato, Yiqun Liu, Miho Sugimoto, Qinglei Wang, and Naoki Orii. Overview of the NTCIR-9 IMine task. In *Proceedings of the 9th NTCIR Conference, NTCIR '09*, Tokyo, Japan, 2011a.
- Yang Song, Dengyong Zhou, and Li-wei He. Post-ranking query suggestion by diversifying search results. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '11*, pages 815–824, New York, NY, USA, 2011b. ACM.
- David Sontag, Kevyn Collins-Thompson, Paul N. Bennett, Ryen W. White, Susan Dumais, and Bodo Billerbeck. Probabilistic models for personalizing web search. In *Proceedings of the 5th ACM International Conference on Web Search and Data Mining, WSDM '12*, pages 433–442, New York, NY, USA, 2012. ACM.
- Alisa Strizhevskaya, Alexey Baytin, Irina Galinskaya, and Pavel Serdyukov. Actualization of query suggestions using query logs. In *Proceedings of the 21st International World Wide Web Conference, WWW '12*, pages 611–612, New York, NY, USA, 2012. ACM.
- Xu Sun, Jianfeng Gao, Daniel Micol, and Chris Quirk. Learning phrase-based spelling error models from clickthrough data. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, ACL '10*, pages 266–274, Stroudsburg, PA, USA, 2010. Association for Computational Linguistics.
- Xu Sun, Anshumali Shrivastava, and Ping Li. Fast multi-task learning for query spelling correction. In *Proceedings of the 21st ACM Conference on Information and Knowledge Management, CIKM '12*, pages 285–294, New York, NY, USA, 2012. ACM.
- Hisami Suzuki and Jianfeng Gao. A unified approach to transliteration-based text input with online spelling correction. In *Proceedings of the 2012 Conference on Empirical Methods in Natural Language Processing, EMNLP '12*, pages 609–618, Stroudsburg, PA, USA, 2012. Association for Computational Linguistics.
- Andrew Swiffin, John Arnott, J. Adrian Pickering, and Alan Newell. Adaptive and predictive techniques in a communication prosthesis. *Augmentative and Alternative Communication*, 3:181–191, 1987.
- Bin Tan, Xuehua Shen, and ChengXiang Zhai. Mining long-term search history to improve search accuracy. In *Proceedings of the 12th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '06*, pages 718–723, New York, NY, USA, 2006. ACM.

- Jaime Teevan, Daniel J. Liebling, and Gayathri Ravichandran Geetha. Understanding and predicting personal navigation. In *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, WSDM '11*, pages 85–94, New York, NY, USA, 2011. ACM.
- Raghavendra Udupa and Shaishav Kumar. Hashing-based approaches to spelling correction of personal names. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing, EMNLP '10*, pages 1256–1265, Stroudsburg, PA, USA, 2010. Association for Computational Linguistics.
- Yury Ustinovskiy and Pavel Serdyukov. Personalization of web-search using short-term browsing context. In *Proceedings of the 22nd ACM Conference on Information and Knowledge Management, CIKM '13*, pages 1979–1988, New York, NY, USA, 2013. ACM.
- David Vallet and Pablo Castells. Personalized diversification of search results. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '12*, pages 841–850, New York, NY, USA, 2012. ACM.
- Joost van Doorn, Daan Odijk, Diederik Roijers, and Maarten de Rijke. Balancing relevance criteria through multi-objective optimization. In *SIGIR 2016: 39th international ACM SIGIR conference on Research and development in information retrieval*, pages 769–772. ACM, 2016.
- Christophe Van Gysel, Leonid Velikovich, Ian McGraw, and Françoise Beaufays. Garbage modeling for on-device speech recognition. In *Interspeech 2015*, September 2015.
- Gregg Vanderheiden and David Kelso. Comparative analysis of fixed-vocabulary communication acceleration techniques. *Augmentative and Alternative Communication*, 3:196–206, 1987.
- Saúl Vargas, Pablo Castells, and David Vallet. Explicit relevance models in intent-oriented information retrieval diversification. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '12*, pages 75–84, New York, NY, USA, 2012. ACM.
- Ingmar Weber and Carlos Castillo. The demographics of web search. In *Proceedings of the 35rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '10*, pages 523–530, New York, NY, USA, 2010. ACM.
- Stewart Whiting and Joemon M. Jose. Recent and robust query auto-completion. In *Proceedings of the 23rd International World Wide Web Conference, WWW '14*, pages 971–982, New York, NY, USA, 2014. ACM.

- Stewart Whiting, James McMinn, and Joemon M. Jose. Exploring real-time temporal query auto-completion. In *Proceedings of the 12th Dutch-Belgian Information Retrieval workshop*, DIR '13, pages 1–4, Delft, The Netherlands, 2013. Technical University Delft.
- Ian H. Witten and John J. Darragh, editors. *The Reactive Keyboard*. Cambridge University Press, Cambridge, UK, 1992.
- Biao Xiang, Daxin Jiang, Jian Pei, Xiaohui Sun, Enhong Chen, and Hang Li. Context-aware ranking in web search. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '10, pages 451–458, New York, NY, USA, 2010. ACM.
- Chuan Xiao, Jianbin Qin, Wei Wang, Yoshiharu Ishikawa, Koji Tsuda, and Kuni-hiko Sadakane. Efficient error-tolerant query autocompletion. *Proceedings of the VLDB Endowment*, 6(6):373–384, 2013.
- Takehiro Yamamoto, Yiqun Liu, Min Zhang, Zhicheng Dou, Ke Zhou, Ilya Markov, Makoto Kato, Hiroaki Ohshima, and Sumio Fujita. Overview of the NTCIR-12 IMine-2 task. In *Proceedings of the 12th NTCIR Conference*, NTCIR '12, pages 8–26, Tokyo, Japan, 2016.
- Jiang-Ming Yang, Rui Cai, Feng Jing, Shuo Wang, Lei Zhang, and Wei-Ying Ma. Search-based query suggestion. In *Proceedings of the 17th ACM Conference on Information and Knowledge Management*, CIKM '08, pages 1439–1440, New York, NY, USA, 2008. ACM.
- Yisong Yue, Rajan Patel, and Hein Roehrig. Beyond position bias: Examining result attractiveness as a source of presentation bias in clickthrough data. In *Proceedings of the 19th International World Wide Web Conference*, WWW '10, pages 1011–1018, New York, NY, USA, 2010. ACM.
- Aston Zhang, Amit Goyal, Weize Kong, Hongbo Deng, Anlei Dong, Yi Chang, Carl A. Gunter, and Jiawei Han. adaQAC: Adaptive query auto-completion via implicit negative feedback. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '15, pages 143–152, New York, NY, USA, 2015. ACM.
- Yuchen Zhang, Weizhu Chen, Dong Wang, and Qiang Yang. User-click modeling for understanding and predicting search-behavior. In *Proceedings of the 17th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '11, pages 1388–1396, New York, NY, USA, 2011. ACM.
- Xiaofei Zhu, Jiafeng Guo, Xueqi Cheng, Pan Du, and Hua-Wei Shen. A unified framework for recommending diverse and relevant queries. In *Proceedings of the 20th International World Wide Web Conference*, WWW '11, pages 37–46, New York, NY, USA, 2011. ACM.

Xiaofei Zhu, Jiafeng Guo, Xueqi Cheng, and Yanyan Lan. More than relevance: High utility query recommendation by mining users' search behaviors. In *Proceedings of the 21st ACM Conference on Information and Knowledge Management, CIKM '12*, pages 1814–1818, New York, NY, USA, 2012. ACM.

Zeyuan Allen Zhu, Weizhu Chen, Tom Minka, Chenguang Zhu, and Zheng Chen. A novel click model and its applications to online advertising. In *Proceedings of the 3rd ACM International Conference on Web Search and Data Mining, WSDM '10*, pages 321–330, New York, NY, USA, 2010. ACM.