



ICTIR

17

The 3rd ACM International Conference on
the Theory of Information Retrieval

October 1-4 2017 | Amsterdam

Conference handbook

Contents

Sunday at a Glance	3
Conference at a Glance	4
Chairs' Welcome	6
Conference Organization	8
Instructions for Participants	12
Sunday	
Tutorials	13
Workshops	15
Monday	
Keynote Speaker	18
Sessions: Entities and knowledge bases	20
Sessions: Query and document representations	21
Sessions: Evaluation	22
Sessions: Panel Discussion	23
Tuesday	
Keynote Speaker	24
Sessions: Retrieval models	25
Sessions: Social medial	26
Sessions: Poster Boosters	27
Wednesday	
Keynote Speaker	29
Sessions: Interactive and session search	31
Sessions: User modelling	32
Venue Information	33
ICTIR 2018	34
Sponsors and support	35

Sunday at a Glance

Activities

09:00–12:30	Tutorial: Differential Privacy for Information Retrieval
09:00–12:30	Tutorial: Efficiency/Effectiveness Trade-offs in Learning to Rank
14:00–18:00	Tutorial: Bandit Algorithms in Interactive Information Retrieval

09:00–18:00	Workshop: Search-Oriented Conversational AI (SCAI)
14:00–18:00	Workshop: Learning Next Generation Rankers (LEARNER 2017)

Lunch & Coffee

10:30–11:00	Coffee break
12:30–14:00	Lunch
15:30–16:00	Coffee break
18:00	Drinks at CASA

Conference at a Glance

Monday, October 2

- 09:00 Welcome Message
- 09:10 Keynote: **Information Retrieval Meets Game Theory**, Oren Kurland
- 10:10 Coffee Break
- 10:40 Session: **Entities and knowledge bases**
- 12:00 Lunch (at the venue)
- 13:00 Session: **Query and document representations**
- 14:20 Coffee Break
- 14:50 Session: **Evaluation**
- 16:10 Panel session
- 18:30 Welcome Reception at Booking.com

Tuesday, October 3

- 09:00 Keynote: **The Evolution of Computational Advertising**, Suju Rajan
- 10:00 Coffee Break & Posters
- 10:30 Session: **Retrieval models**
- 11:50 Lunch (at the venue)
- 13:00 Session: **Social medial**
- 14:20 Coffee Break & Posters
- 14:50 Poster Booster
- 15:20 Poster Session
- 17:00 Award Announcement
- 19:00 Banquet Dinner at De Kas

Wednesday, October 4

09:00	Keynote: Learning in networks: How to exploit relationships to improve predictions , Jennifer Neville
10:00	Coffee Break
10:30	Session: Interactive and session search
11:50	Lunch (at the venue)
13:00	Business Meeting
14:00	Session: User modelling
15:00	Conference Close

Chairs' Welcome

Welcome to ICTIR 2017, the 7th International Conference on the Theory of Information Retrieval and the 3rd conference with that name to be fully sponsored by the ACM Special Interest Group on Information Retrieval (SIGIR). This year's conference continues its tradition of being the premier forum for presentation of research on theoretical aspects of Information Retrieval (IR) including (a) conceptual papers that explore key concepts, (b) theoretical papers that model concepts and/or relations between concepts, and (c) papers that study theory in experimental or industrial settings. To highlight the increasingly strong connections between Information Retrieval and neighboring disciplines this year's conference explicitly welcomed papers in IR areas that overlap with Human Information Access, Machine Learning, Natural Language Processing and Perception.

We are also happy to welcome you to Amsterdam, The Netherlands. The Netherlands has a number of strong and vibrant IR groups, e.g. at the University of Amsterdam, Twente University, Radboud University Nijmegen, Centrum Wiskunde en Informatica, Delft University of Technology, and University of Leiden and counts about 150 IR researchers, the largest number per capita in worldwide. These groups are successfully engaged in both theoretical and practical work. They are young groups consisting of enthusiastic people, with a track record in high quality publications.

The call for papers attracted submissions from all over the world. The program committee reviewed 97 contributions – 54 long and papers 43 short papers – and accepted 27 long papers and 25 short papers, all of which will be presented during the conference. We received three excellent tutorial proposals and two workshop proposals, and we hope those will aid in laying down foundational ideas in information retrieval research. We thank everyone who helped us organize the conference: short paper chairs Katja Hofmann and Christina Lioma, workshop chair Grace Hui Yang, tutorial chair Josiane Mothe, core IR track program chairs Peter Bruza and Lynda Tamine-Lechani, IR & Human Information Access track chair Diane Kelly, IR & Machine Learning track chairs Qiaozhu Mei and Alessandro Sordoni, IR & Natural Language Processing track chair Alessandro Moschetti, IR & Perception track chairs Maria Eskevich and Stefan Rueger, and the members of the program committee who worked hard to review papers and provide feedback for authors. We also thank our sponsors and supporters: ACM SIGIR, Bloomberg, Booking.com, Facebook,

Gemeente Amsterdam, Google, Microsoft, NWO, SIKS, University of Amsterdam, Werkgemeenschap Informatiewetenschap. Finally, we thank the authors for providing superb technical content, the workshop organizers, the tutorial lecturers and our keynote speakers Oren Kurland, Jennifer Neville and Suju Rajan.

We hope that you will find the conference program interesting and thought-provoking and that the conference will provide you with a valuable opportunity to share ideas with other researchers and practitioners from institutions around the world.

Enjoy the conference!

**Jaap Kamps, Evangelos Kanoulas &
Maarten de Rijke**

ICTIR'17 General Chairs

University of Amsterdam

**Hui Fang &
Emine Yilmaz**

ICTIR'17 Program Chairs

University of Delaware &
University College London

Conference Organization

General Chairs

Jaap Kamps, University of Amsterdam, Netherlands
Evangelos Kanoulas, University of Amsterdam, Netherlands
Maarten de Rijke, University of Amsterdam, Netherlands

Program Chair

Hui Fang, University of Delaware, USA
Emine Yilmaz, University College London, UK

Short Papers Chair

Katya Hofmann, Microsoft Research, UK
Christina Lioma, University of Copenhagen

Tutorial Chair

Josiane Mothe, IRIT, France

Workshop Chair

Grace Hui Yang, Georgetown University, USA

IR Track Chair

Peter Bruza, Queensland University of Technology, Australia
Lynda Tamine-Lechani, IRIT, France

IR & Human Information Access Track Chair

Diane Kelly, University of Tennessee, USA

IR & Machine Learning Track Chair

Qiaozhu Mei, University of Michigan, USA
Alessandro Sordoni, University of Michigan, USA

IR & Perception Track Chair

Maria Eskevich, Radboud University, Netherlands
Stefan Rueger, The Open University, UK

IR & Natural Language Processing Chair

Alessandro Moschetti, Hamad bin Khalifa University, Qatar

Publicity Chair

Julia Kiseleva, University of Amsterdam, Netherlands

Proceeding Chair

Hosein Azarbonyad, University of Amsterdam, Netherlands

Conference Handbook Chair

Christophe Van Gysel, University of Amsterdam, Netherlands

Dan Li, University of Amsterdam, Netherlands

Chang Li, University of Amsterdam, Netherlands

Program Committee

Enrique Amigó, UNED, Spain

Jaime Arguello, University of North Carolina at Chapel Hill, USA

Leif Azzopardi, University of Strathclyde, United Kingdom

Phil Bachman, Microsoft, Canada

Alberto Barrón-Cedeño, Qatar Computing Research Institute, Qatar

Gaurav Baruah, University of Waterloo, Canada

Pavel Braslavski, Ural Federal University, Russia

Jamie Callan, Carnegie Mellon University, USA

Pablo Castells, Universidad Autónoma de Madrid, Spain

Chang Chia-Hui, National Central University, Taiwan

Kevyn Collins-Thompson, University of Michigan, USA

Ingemar Cox, University College London, United Kingdom

Matt Crane, University of Waterloo, Canada

Fabio Crestani, University of Lugano, Switzerland

Danilo Croce, University of Roma Tor Vergata, Italy

Bruce Croft, University of Massachusetts Amherst, USA

Ronan Cummins, University of Cambridge, United Kingdom

Jeff Dalton, University of Massachusetts Amherst, USA

Kareem Darwish, Qatar Computing Research Institute, Qatar

Arjen de Vries, Radboud University, Netherlands

Gianluca Demartini, University of Sheffield, United Kingdom

Tamer Elsayed, Qatar University, Qatar

Ralph Ewerth, Leibniz Universität Hannover, Germany
James Fan, HelloVera.ai, USA
Nicola Ferro, University of Padua, Italy
Simone Filice, University of Roma Tor Vergata, Italy
Ingo Frommholz, University of Bedfordshire, United Kingdom
Norbert Fuhr, University of Duisburg-Essen, Germany
Shlomo Geva, Queensland University of Technology, Australia
Julio Gonzalo, UNED, Spain
Martin Halvey, University of Strathclyde, United Kingdom
Morgan Harvey, Northumbria University, United Kingdom
Ben He, University of Chinese Academy of Sciences, China
Jiyin He, CWI, Netherlands
Djoerd Hiemstra, University of Twente, Netherlands
Jimmy Huang, York University, Canada
Mohit Iyyer, University of Maryland, USA
Jing Jiang, Singapore Management University, Singapore
Jussi Karlgren, Gavagai & KTH, Sweden
Bevan Koopman, CSIRO, Australia
Wai Lam, The Chinese University of Hong Kong, Hong Kong
Birger Larsen, Aalborg University Copenhagen, Denmark
Cheng Li, University of Michigan, USA
Hang Li, Huawei Technologies, Hong Kong
Joo-Hwee Lim, Nanyang Technological University, Singapore
Chang Liu, Peking University, China
Haiming Liu, University of Bedfordshire, United Kingdom
Yiqun Liu, Tsinghua University, China
Joao Magalhaes, Universidade Nova de Lisboa, Portugal
Ilya Markov, University of Amsterdam, Netherlands
Massimo Melucci, University of Padua, Italy
Bhaskar Mitra, Microsoft, United Kingdom
Alistair Moffat, The University of Melbourne, Australia
Boughanem Mohand, IRIT University Paul Sabatier Toulouse, France
Yashar Moshfeghi, University of Strathclyde, United Kingdom
Jayanta Mukherjee, IIT Kharagpur, India
Jian-Yun Nie, Université de Montréal, Canada
Javier Parapar, University of A Coruña, Spain
Benjamin Piwowarski, CNRS / University Pierre et Marie Curie, France
Simone Paolo Ponzetto, University of Mannheim, Germany
Tao Qin, Microsoft Research, China
Filip Radlinski, Google, United Kingdom

Jinfeng Rao, University of Maryland, USA
Salvatore Romeo, Qatar Computing Research Institute, Qatar
Ian Ruthven, University of Strathclyde, United Kingdom
Rodrygo Santos, Universidade Federal de Minas Gerais, Brazil
Fabrizio Sebastiani, Qatar Computing Research Institute, Qatar
Aliaksei Severyn, Google, Switzerland
Chirag Shah, Rutgers University, USA
Yangyang Shi, Microsoft, USA
Laurianne Sitbon, Queensland University of Technology, Australia
Ian Soboroff, NIST, USA
Laure Soulier, Sorbonne Universités UPMC-LIP6, France
Paul Thomas, Microsoft, Australia
Gabriele Tolomei, University of Padua, Italy
Kentaro Torisawa, NICT, Japan
Andrew Trotman, University of Otago, New Zealand
Theodora Tsikrika, Information Technologies Institute, CERTH, Greece
Ellen Voorhees, NIST, USA
Stefanos Vrochidis, Information Technologies Institute, Greece
V.G.Vinod Vydiswaran, University of Michigan, USA
Bin Wang, Chinese Academy of Sciences, China
Jun Wang, University College London, United Kingdom
Xuanhui Wang, Google, USA
Yue Wang, University of Michigan, USA
Max Wilson, University of Nottingham, United Kingdom
Rui Yan, Peking University, China
Liu Yang, University of Massachusetts Amherst, USA
Guido Zuccon, Queensland University of Technology, Australia

Additional reviewers

Maram Barifah
Seyed Ali Bahrainian
Anastasia Giachanou
Maram Hasanain
Mucahid Kutlu
Mengwen Liu
Maria Maistro
Ida Mele
Matthew Mitsui
Mohammad Alian Nejadi
Thibaut Thonet
Nikos Voskarides

Instructions for Participants

Talks

Full paper presentations are 20 minutes long. **Short papers** will be presented during the poster booster session (1 minute). We strongly encourage authors of short papers to use the time as an introduction to your work rather than a technical deep-dive. You can cover technical details in interactive discussion at the poster sessions.

Posters

We have one poster session on Tuesday: **14:50 –17:00**. Please be present at your poster during the poster session. Your poster will be displayed on corridor windows of the conference venue; hence there are no restrictions regarding the size of your poster, although we would recommend an A3 portrait.

Poster setup and takedown

Posters can stay at the conference windows during the entire conference. Time for poster setup is Tuesday morning. Poster takedown happens after the last session of the conference (but you are also welcome to remove it by the end of the poster session).

Differential Privacy for Information Retrieval

*Grace Hui Yang (Georgetown University),
Sicong Zhang (Georgetown University)*

09:00–12:30

Information Retrieval (IR) research has extensively utilized personalization to advance its state-of-the-art. In this process, many IR algorithms and applications require the use of users' personal information, contextual information and other sensitive and private information. However, while IR researchers are making progress, there is always a concern of how not to violate the users' privacy. Sometimes, the concern becomes so overwhelming that the IR research has to be stopped to avoid leaking of users' privacy. The good news is that recently increasing attentions have been paid on the joint field of privacy and IR – privacy-preserving IR. As part of the effort, this tutorial offers an introduction to differential privacy (DP), one of the most advanced techniques in privacy research, and provides necessary set of theoretical knowledge for applying privacy techniques in IR. Differential privacy is a technique that provides strong privacy guarantees for data protection. Theoretically, it aims to maximize the data utility in statistical datasets while minimizing the risk of exposing individual data entries to any adversary. Differential privacy has been applied across of a wide range of applications in database, data mining, and IR. This tutorial aims to lay a theoretical foundation of DP and how it can be applied to IR.

Bandit Algorithms in Interactive Information Retrieval

*Dorota Glowacka (Department of Computer Science
University of Helsinki)*

14:00–18:00

The multi-armed bandit problem models an agent that simultaneously attempts to acquire new knowledge (exploration) and optimize his decisions based on existing knowledge (exploitation). The agent attempts to balance these competing tasks in order to maximize his total value over the period of time considered. There are many practical applications of the bandit model, such as clinical trials, adaptive routing or portfolio design. Over the last decade there has been an increased interest in the development of new bandit algorithms for specific problems in information, such as diverse document ranking, news recommendation or ranker evaluation. The aim of this tutorial is to provide an overview of the various applications of bandit algorithms in information retrieval as well as issues related to their practical deployment and performance in real-life systems/applications.

Efficiency/Effectiveness Trade-offs in Learning to Rank

*Claudio Lucchese (ISTI-CNR),
Franco Maria Nardini (ISTI-CNR)*

9:00–12:30

Tutorial on the efficiency/effectiveness tradeoffs in Learning to Rank. In the last years, learning to rank (LtR) had a significant influence on several tasks in the Information Retrieval field, with large research efforts coming both from the academia and the industry. Indeed, efficiency requirements must be fulfilled in order to make an effective research product deployable within an industrial environment. The evaluation of a model can be too expensive due to its size, the features used and several other factors. This tutorial discusses the recent solutions that allow to build an effective ranking model that satisfies temporal budget constraints at evaluation time. For more information please visit the tutorial website, <http://learningtorank.isti.cnr.it/>.

Learning Next Generation Rankers (LEARNER 2017)

*Nicola Ferro (Dept. of Information Engineering,
University of Padua, Italy),*

Claudio Lucchese (ISTI-CNR, Pisa, Italy),

*Maria Maistro (Dept. of Information Engineering,
University of Padua, Italy),*

Raffaele Perego (ISTI-CNR, Pisa, Italy)

14:00–18:00

The aim of LEARNER@ICTIR2017 is to investigate new solutions for LtR. In details, we identify some research areas related to LtR which are of actual interest and which have not been fully explored yet. We solicit the submission of position papers on novel LtR algorithms, on evaluation of LtR algorithms, on dataset creation and curation, and on domain specific applications of LtR. LEARNER@ICTIR2017 will be a gathering of academic people interested in IR, ML and related application areas. We believe that the proposed workshop is relevant to ICTIR since we look for novel contributions to LtR which focus on foundational and conceptual aspects, which need to be properly framed and modelled.

Search-Oriented Conversational AI (SCAI)

*Mikhail Burtsev (MIPT, Russia),
Aleksandr Chuklin (Google Research Europe, Switzerland),
Julia Kiseleva (University of Amsterdam, Netherlands),
Alexey Borisov (Yandex & University of Amsterdam)*

09:00–18:00

The aim of SCAI@ICTIR2017 is to bring together IR and AI communities to instigate future direction of search-oriented conversational systems. We identified the number of research areas related to conversational AI which is of actual interest to both communities and which have not been fully explored yet. We think it's beneficial to exchange our visions. We solicit the paper submissions and more importantly proposals for panel discussions where researchers can exchange opinions and experiences. We believe that the proposed workshop is relevant to ICTIR since we look for novel contributions to search-oriented conversational systems which are a new and promising area.

Information Retrieval Meets Game Theory

Oren Kurland

Technion – Israel Institute of Technology

09:10–10:10

Chair: Hui Fang (University of Delaware)

Abstract

In competitive search settings such as the Web, authors of documents may have an incentive to have their documents highly ranked for certain queries. This can drive corpus dynamics as documents may be manipulated in response to induced rankings (e.g., by applying search engine optimization). Such post-ranking corpus effects are not directly modelled in ad hoc retrieval models and, more generally, are not accounted for by the formal foundations of retrieval paradigms. In this talk I will discuss how (algorithmic) game theory can be used to analyze some aspects of the competitive search setting. I will first discuss the probability ranking principle (PRP) which is the theoretical underpinning of most ad hoc retrieval methods. As it turns out, the PRP is sub-optimal in competitive settings. In addition, I will discuss some initial theoretical and empirical results regarding the strategic behavior of document authors in competitive retrieval settings, specifically with respect to the foundations of classical ad hoc retrieval models. I will then discuss future directions.

Bio

Oren Kurland is an Associate Professor at the Technion – Israel Institute of Technology. He holds a Ph.D. in Computer Science from Cornell University. Oren's main research focus is information retrieval. He serves on the editorial board of the Information Retrieval Journal, and has also served on the editorial boards of the Journal of Artificial Intelligence Research and Information Processing and Management Journal. Oren served as the first chair of the ACM SIGIR ICTIR steering committee. He has also served as program committee co-chair of the ICTIR and SPIRE conferences. Oren was awarded best-paper

Monday Keynote Speaker

honorable mention awards in the SIGIR 2013 and ECIR 2016 conferences. He has received faculty research awards from IBM, Google and Yahoo.

Monday Sessions

Entities and knowledge bases

Chair: Shane Culpepper

Talks (10:40–12:00)

- 10:40 **Structural Regularities in Text-based Entity Vector Spaces**
Christophe Van Gysel, Maarten de Rijke and Evangelos Kanoulas
- 11:00 **Knowledge Questions from Knowledge Graphs**
Dominic Seyler, Mohamed Yahya and Klaus Berberich
- 11:20 **A Deep Neural Information Retrieval Model Enhanced by a Knowledge Resource Driven Representation of Documents**
Dominic Seyler, Mohamed Yahya and Klaus Berberich
- 11:40 **On Type-Aware Entity Retrieval**
Dario Garigliotti and Krisztian Balog
- 12:00 **End of session**

Monday Sessions

Query and document representations

Chair: Djoerd Hiemstra

Talks (13:00–14:20)

- 13:00 **Enhanced Mean Retrieval Score Estimation for Query Performance Prediction**
Haggai Roitman, Shai Erera, Oren Sar Shalom and Bar Weiner
- 13:20 **Can Short Queries Be Even Shorter?**
Peilin Yang and Hui Fang
- 13:40 **Deriving Differentially Private Session Logs for Query Suggestion**
Sicong Zhang and Grace Hui Yang
- 14:00 **Graph-based semi-supervised learning for text classification**
Natalie Widmann and Suzan Verberne
- 14:20 **End of session**

Evaluation

Chair: Fabio Crestani

Talks (14:50–18:10)

- 14:50 **Are IR evaluation measures on an interval scale?**
Marco Ferrante, Nicola Ferro and Silvia Pontarollo
- 15:10 **Considering Assessor Agreement in IR Evaluation**
Eddy Maddalena, Kevin Roitero, Gianluca Demartini and Stefano Mizzaro
- 15:30 **Dealing with Incomplete Judgments in Cascade Measures**
Kai Hui, Klaus Berberich and Ida Mele
- 15:50 **Evaluation Measures for Relevance and Credibility in Ranked Lists**
Christina Lioma, Jakob Grue Simonsen and Birger Larsen
Fabio Crestani
- 16:10 **End of session**

Panel Discussion
Chair: Arjen de Vries

The end of IR as we know it? (16:10–17:10)

Arjen de Vries

Radboud Universiteit

Jimmy Lin

University of Waterloo

Grace Hui Yang

Georgetown University

Norbert Fuhr

University of Duisburg-Essen

Fiana Raiber

Yahoo Research

Leif Azzopardi

University of Strathclyde

The Evolution of Computational Advertising

Suju Rajan

Criteo Research (France)

09:00–10:00

Chair: Lynda Tamine-Lechani (Paul Sabatier University)

Abstract

Machine learning literature on computational advertising typically tends to focus on the simplistic CTR prediction problem which while being relevant is the tip of the iceberg in terms of the research challenges in the field. There have been several recent efforts, shaped by the realities of a complex ad ecosystem, to develop models that try to better encapsulate the journey of an ad from its impression to possibly leading to a purchase. In this talk, we will highlight the recent research challenges in the field of computational advertising & how it is evolving to incorporate ideas from areas such as reinforcement learning, econometrics, deep learning & large-scale recommender systems.

Bio

Suju Rajan is the VP, Head of Research at Criteo. At Criteo, her team works on all aspects of performance driven computational advertising, including, real-time bidding, large-scale recommendation systems, auction theory, reinforcement learning, online experimentation, metrics and scalable optimization methods. Prior to Criteo, she was the Director of the Personalization Sciences at Yahoo Research where her team worked on personalized recommendations for several Yahoo products. She received her PhD from the University of Texas at Austin.

Tuesday Sessions

Retrieval models

Chair: Norbert Fuhr

Talks (10:30–11:50)

- 10:30 **Improved Query Topic Models Using Pseudo-Relevant Polya Document Models**
Ronan Cummins
- 10:50 **On the Additivity and Weak Baselines for Search Result Diversification Research**
Mehmet Akcay, Ismail Sengor Altingovde, Craig Macdonald and Iadh Ounis
- 11:10 **Kullback-Leibler Divergence Revisited**
Fiana Raiber and Oren Kurland
- 11:30 **Text Retrieval based on Least Information Measurement**
Weimao Ke
- 11:50 **End of session**

TUESDAY

Tuesday Sessions

Social medial

Chair: Ismail Sengor Altingovde

Talks (14:00–14:30)

- 13:00 **Mining Temporal Statistics of Query Terms for Searching Social Media Posts**
Jinfeng Rao, Ferhan Ture, Xing Niu and Jimmy Lin
- 13:20 **Modeling Controversy within Populations**
Myungha Jang, Shiri Dori-Hacohen and James Allan
- 13:40 **Recommendation with Social Relationships via Deep Learning**
Dimitrios Rafailidis and Fabio Crestani
- 14:00 **Characterizing and Predicting Supply-side Engagement on Crowd-contributed Video Sharing Platforms**
Rishabh Mehrotra and Prasanta Bhattacharya
- 14:20 **End of session**

Poster Boosters

Chair: Krisztian Balog

Talks (14:50–15:20)

- Term-Mouse-Fixations as an Additional Indicator for Topical User Interests in Domain-Specific Search by Daniel Hienert and Dagmar Kern
- On Search Powered Navigation by Mostafa Dehghani, Glorianna Jagfeld, Hosein Azarbonyad, Alex Olieman, Jaap Kamps and Maarten Marx
- Mobile Vertical Ranking based on Preference Graphs by Yuta Kadotami, Yasuaki Yoshida, Sumio Fujita and Tetsuya Sakai
- Personalised Search Time Prediction using Markov Chains by Tuan Vu Tran, David Maxwell, Norbert Fuhr and Leif Azzopardi
- An Exploration of Serverless Architectures for Information Retrieval by Matt Crane and Jimmy Lin
- Robust Standard Deviation Estimation for Query Performance Prediction by Haggai Roitman, Shai Erera and Bar Weiner
- The Pareto Frontier of Utility Models as a Framework for Evaluating Push Notification Systems by Gaurav Baruah and Jimmy Lin
- Advanced Hidden Markov Models for Recognizing Search Phases by Sebastian Dungs and Norbert Fuhr
- Uncovering Like-minded Political Communities on Twitter by Ophélie Fraiser, Guillaume Cabanac, Yoann Pitarch, Romaric Besançon and Mohand Boughanem
- Quantization in Append-Only Collections by Salman Mohammed, Matt Crane and Jimmy Lin
- Upper Bound Approximations for BlockMaxWand by Nicola Tonellotto and Craig Macdonald
- Merge-Tie-Judge: Low-Cost Preference Judgments with Ties by Kai Hui and Klaus Berberich

- Evaluating and Analyzing Click Simulations in Web Search by Stepan Malkevich, Ilya Markov, Maarten de Rijke and Elena Mikhailova
- An Initial Investigation of Query Expansion Bias by Colin Wilkie and Leif Azzopardi
- Towards Learning Reward Functions from User Interactions by Ziming Li, Julia Kiseleva, Maarten de Rijke and Artem Grotov
- Benchmark for Complex Answer Retrieval by Federico Nanni, Bhaskar Mitra, Matt Magnusson and Laura Dietz
- Detecting Seasonal Queries Using Time Series and Content Features by Behrooz Mansouri, Mohammad Sadegh Zahedi, Maseud Rahgozar and Ricardo Campos
- A Contextual Bandit Approach to Dynamic Search by Angela Yang and Grace Hui Yang
- Modeling Information Flow in Dynamic Information Retrieval by Felipe Moraes, MÃ¡rio Alvim and Rodrygo Santos
- On the Effectiveness of Bayesian Network-based Models for Document Ranking by Xing Tan, Fanghong Jian and Jimmy Huang
- On Fine-Grained Geo-Localisation of Tweets by Jorge David Gonzalez Paule, Yashar Moshfeghi, Joemon M. Jose and Piyushimita Thakuriah
- The Treatment of Ties in AP Correlation by Julian Urbano and MÃ¡nica Marrero
- Enhanced Probabilistic Classify and Count Methods for Multi-Label Text Quantification by Roy Levin and Haggai Roitman
- Retrieving Compositional Documents Using Position-Sensitive Word Mover's Distance by Martin Trapp, Marcin Skowron and Dietmar Schabus
- Emotion Detection from Text via Ensemble Classification Using Word Embeddings by Jonathan Herzig, Michal Shmueli-Scheuer and David Konopnicki

Learning in networks: How to exploit relationships to improve predictions

Jennifer Neville

Purdue University

09:00–10:00

Chair: Julia Kiseleva (University of Amsterdam)

Abstract

The popularity of social networks and social media has increased the amount of information available about users' behavior online – including current activities, and interactions with followers, friends, and family. This rich relational information can be used to improve predictions even when individual data is sparse, since the characteristics of friends are often correlated. Although this type of network data offer several opportunities to improve predictions about users, the characteristics of online social network data also present a number of challenges to accurately incorporate the network information into machine learning systems. This talk will outline some of the algorithmic and statistical challenges that arise due to partially-observed, large-scale networks, and describe methods for semi-supervised learning, latent-variable modeling, and active sampling to address the challenges.

This talk will describe some of the tasks and sub-areas of NLP that have experienced fundamental advances in the state of the art over past few years, focusing on what these advances mean for understanding and extracting information from online text.

Bio

Jennifer Neville is the Miller Family Chair Associate Professor of Computer Science and Statistics at Purdue University. She received her PhD from the University of Massachusetts Amherst in 2006. She is currently an elected member of the AAAI Executive Council and she was recently PC chair of the 9th ACM International Conference on Web Search and Data. In 2012, she was awarded an NSF Career Award,

Wednesday Keynote Speaker

in 2008 she was chosen by IEEE as one of “AI’s 10 to watch”, and in 2007 was selected as a member of the DARPA Computer Science Study Group. Her work, which includes more than 100 peer-reviewed publications with over 5000 citations, focuses on developing data mining and machine learning techniques for complex relational and network domains, including social, information, and physical networks.

Wednesday Sessions

Interactive and session search

Chair: Haggai Roitman

Talks (10:30–11:50)

- 10:30 **Towards the Next Generation of Personal Assistants:
Systems that Know When You Forget**
Seyed Ali Bahrainian and Fabio Crestani
- 10:50 **On Effective Dynamic Search in Specialized Domains**
Felipe Moraes, Rodrygo Santos and Nivio Ziviani
- 11:10 **Investigating Per Topic Upper Bound for Session
Search Evaluation**
Zhiwen Tang and Grace Hui Yang
- 11:30 **Information Retrieval Evaluation as Search Simula-
tion: A General Formal Framework for IR Evaluation**
Yinan Zhang, Xueqing Liu and Chengxiang Zhai
- 11:50 **End of session**

Wednesday Sessions

User modelling
Chair: Nicola Ferro

Talks (14:00–15:00)

- 14:00 **A User Re-Modeling Approach to Item Recommendation using Complex Usage Data**
Oren Sar Shalom, Haggai Roitman, Yishay Mansour and Amihod Amir
- 14:20 **The Positive and Negative Influence of Search Results on People's Decisions about the Efficacy of Medical Treatments**
Frances Pogacar, Amira Ghenai, Mark Smucker and Charles Clarke
- 14:40 **Personalized Navigation and Random Walk on the Complex Heterogeneous Graph**
Xiaozhong Liu, Yingying Yu, Zhuoren Jiang, Chun Guo and Scott Jensen
- 15:00 **End of session**

Venue Information

ICTIR 2017 will be held at Hotel Casa, in Amsterdam, Netherlands. The Hotel Casa is located on the south of Amsterdam near Amstel. *The registration desk* will be open the whole day during the conference. *Drinks* will be offered on Sunday at the Wine Bar. *Reception* will be held on Monday at Booking.com. Entry only with ICTIR badge at the Herengracht 597 entrance (the one with the reception desk). *Conference Banquet* will be at De Kas restaurant on Tuesday evening.

Local Information

Q&A about Amsterdam

Should I tip & how much? In the Netherlands, prices include taxes and service, but leaving a small tip (5-10% of the total bill including taxes) is very common.

Why is my credit card not accepted? The Netherlands uses its own debit card system (called PIN). Many restaurants and larger shops do accept VISA and MasterCard, but check beforehand.

What should I buy my family? Stroopwafels (if you love them), drop (if you don't), cheese (Old Amsterdam is nice) or tulips. Get any of the first three from a supermarket like Albert Heijn. For tulips, there is a flower market near the Muntplein (very close to the conference venue). In a pinch on your way to the airport, check out the IAmsterdam Store in the back of Amsterdam Central Station.

Getting Around

Tram, bus, metro plan your trip with Google maps; buy one-hour tickets for 2.90 Euro at stations and in trams (but not in buses)

Taxi www.tcataxi.nl/en (+31 20 7 777 777) or use Uber

City portal for visitors www.iamsterdam.com

ICTIR 2018

ICTIR 2018

The Fourth ACM SIGIR International Conference on the Theory of Information Retrieval

September 24-28, Tianjin, China

<http://ictir2018.org>

SPONSORS

ACM SIGIR
Tianjin University
Microsoft Research Asia
Chinese Information Processing Society of China



IMPORTANT DATES (Tentative)

- Long paper submission deadline: April 2018
- Short paper submission deadline: April 2018
- Tutorial submission deadline: April 2018
- Workshop submission deadline: April 2018
- Notifications of acceptance: May/June 2018

ORGANISERS

- **General Co-chairs:**
 - Dawei Song (*The Open University, UK & Tianjin University, China*)
 - Tie-Yan Liu (*Microsoft Research Asia, China*)
 - Le Sun (*Chinese Information Processing Society of China*)
- **Programme Committee Co-chairs (Long papers):**
 - Peter Bruza (*Queensland University of Technology*)
 - Massimo Melucci (*University of Padua, Italy*)
- **Programme Committee Co-chairs (Short papers):**
 - Fabrizio Sebastiani (*Institute for Science and Technologies of Information, Italy*)
 - Grace Hui Yang (*Georgetown University, USA*)

Venue:

Tianjin Yingbin Hotel



Sponsors and support



Part of:

