On the Volatility of Commercial Search Engines and its Impact on Information Retrieval Research

Jimmy
Queensland University of Technology
Brisbane, Australia
University of Surabaya (UBAYA)
Surabaya, Indonesia
jimmy@hdr.qut.edu.au

Guido Zuccon
Queensland University of Technology
Brisbane, Australia
g.zuccon@qut.edu.au

Gianluca Demartini
University of Queensland
Brisbane, Australia
g.demartini@uq.edu.au

ABSTRACT

We studied the volatility of commercial search engines and reflected on its impact on research that uses them as basis of algorithmical techniques or for user studies. Search engine volatility refers to the fact that a query posed to a search engine at two different points in time returns different documents.

By comparing search results retrieved every 2 days over a period of 64 days, we found that the considered commercial search engine API consistently presented volatile search results: it both retrieved new documents, and it ranked documents previously retrieved at different ranks throughout time. Moreover, not only results are volatile: we also found that the effectiveness of the search engine in answering a query is volatile. Our findings reaffirmed that results from commercial search engines are volatile and that care should be taken when using these as basis for researching new information retrieval techniques or performing user studies.

ACM Reference Format:
Jimmy, Guido Zuccon, and Gianluca Demartini. 2018. On the Volatility of Commercial Search Engines and its Impact on Information Retrieval Research. In *SIGIR '18: The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, July 8–12, 2018, Ann Arbor, MI, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3209978.3210088

1 INTRODUCTION

On a number of occasions, information retrieval researchers have used commercial search engines and associated APIs to assist with the research of new algorithms and techniques (type A: algorithmical use), or to investigate user search behaviour (type U: user study use). Examples of this practice include, among others: Cilibrasi and Vitanyi [5] defined a word similarity function based on the number of search results retrieved by Google (A); Symonds et al. [10] used Google to perform a first round of retrieval to inform query expansion (A); Maxwell et al. [8] used Bing to retrieve documents and snippets within a user study that explored user behaviour with respect to snippet length and informativeness trade-off (U).

To help understand the extent of this practice, we systematically surveyed the literature published in the ACM SIGIR conference between 2006 and 2016 (a total of 2,138 full and short papers)[1]. We found that 158 contributions (7.4%) used commercial search engines in their experiments[2].

Commercial search engines are however often *volatile*: both the results retrieved and their rankings often differ given two points in time. There are multiple reasons for this volatility. On the one hand, volatility may be due to index updates and fresher information being indexed by the search engines [3, 4]. On the other hand, volatility may be due to operational reasons such as index replication, sharding and routing. Differences may also be due to updates to the ranking function used by the search engines. McCown and Nelson [9] also found that commercial search engines' API and their web interfaces often access different indexes: they argued that the API may access a smaller index than the web interface. Our study focuses on measuring the volatility of search results returned by a search engine's API.

Commercial search engine volatility has been investigated by Altingovde et al. [2] who had experimented using a set of 630,000 queries and found that only 10.7% of top 10 results found in 2007 remained as top 10 results in 2010. Bai and Junqueira studied volatility in Yahoo! and reported that of 1.4M search results analysed over 3 weeks, ≈ 35% new URLs were added, ≈ 1% had modified content, and ≈ 0.06% were deleted [3]. Specific to the rate of URLs with modified content, Adar et al. [1] found that 34% of URLs in their 5 weeks study had no change, while the remaining changed on average every 123 hours (average Dice coefficient: 0.794). These changes may had impacted the search quality over time: however that study did not explicitly measure the relevancy of the search results and whether volatility impacted on search engine effectiveness. On the contrary, we also measure changes in results relevance and search engine effectiveness over time.

Search engine volatility may be a problem when commercial search engines are used by researchers as part of their methods or user studies. In other words: if an algorithm or technique is based on the use of a volatile search service, differences in search results and rankings may vary the effectiveness of the method, or render the replication of the experiments impossible. Similarly, if a user study relies on a commercial search engine to investigate user behaviour, volatility may be a confounding factor affecting effectiveness, especially if the user studies are carried out over a

SIGIR '18, July 8–12, 2018, Ann Arbor, MI, USA

ACM ISBN 978-1-4503-5657-2/18/07…$15.00
https://doi.org/10.1145/3209978.3210088

[1]Note that this practice is well utilised also outside of the SIGIR literature, as demonstrated by the examples cited above that were not published in SIGIR.
[2]Data available at http://ielab.io/se-volatility

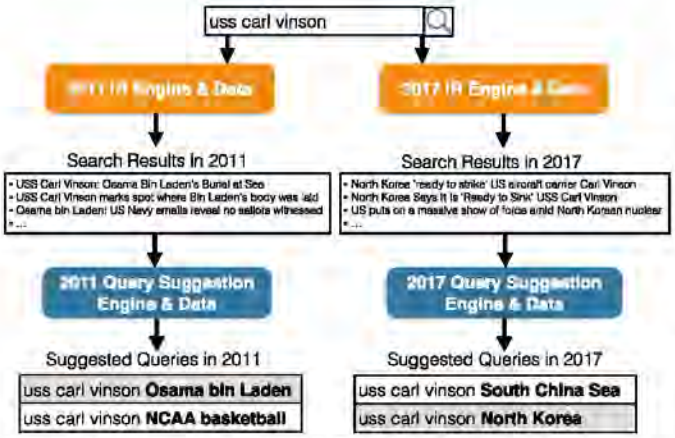

Figure 1: Example of search result volatility and its impact on an hypothetical query expansion techniques that exploits the retrieved results.

period of time, rather than all being run concurrently. Yet, this aspect is often ignored when analysing the results, e.g., volatility is not considered as factor within an ANOVA analysis of results.

To further exemplify how search engine volatility may affect information retrieval research that relies on such commercial services, consider the case of a user study relying on the Bing search APIs and investigating the capabilities of users in selecting query suggestions automatically generated by techniques that exploit the search results obtained from the initial user's query. In such case, two components of the experimental methodology rely on results from the commercial APIs: (i) the query suggestion mechanism, and (ii) the results that are retrieved (and evaluated for effectiveness) in response to the user's query and the selected query suggestion. Figure 1 shows an example of a query submitted at two different times. At each time, the query retrieved two (sensibly) different sets of results and thus returned different query suggestions.

In this study, we seek to answer the following research questions:

- RQ1: What amount of volatility do commercial search engine APIs present?
- RQ2: How does the volatility of commercial search engines affect information retrieval research?

To answer RQ1, we periodically used a search engine API to retrieve results for a large set of queries which had no specific temporal intent or seasonality effect. We then studied how results changed over time. To answer RQ2 we assessed the relevance of the top search results we collected over time and we analysed the change in search engine effectiveness over time. Details of the methods used in this study are described next.

2 METHODS

To answer RQ1, we acquired the queries used in the TREC 2013 and 2014 Web Track (100 queries in total). While these queries had no explicit temporal nor seasonal intent, a small number may have been influenced by temporal issues. For example, in our experiments, query 202: "uss carl vinson" was affected by the US decision of deploying the aircraft carrier within strike range of North Korea in early January 2018. We further acquired a set of 300 queries from the CLEF 2016 eHealth IR collection. Of the 300 queries, we removed query 129005 due to a problem with quotation mark characters in the query. These queries related to consumer health search intents, and were unlikely to be affected by temporal or seasonal intents.

We used the Bing Search API [3] to retrieve a maximum of 50 web results in answer to the query sets, setting `English US` as the market and with `safe search` turned off. We performed retrieval every two days from 29/11/2017 to 31/01/2018 with exceptions of 13, 27, and 29 December 2017 where the retrieval process was not triggered due to technical problems. Hence, in total we collected 30 data samples for each query set. We used the data samples to investigate the volatility of search API by counting the number of new URLs between search results from different retrieval dates pairs. We then further investigated whether differences in ranking for an URL were also found over time. To determine whether IP address or browser cookies may have affected the results, we repeated some of the crawls using an additional server located in another Australian state; crawls were started at the same time. A 99.98% match was found between the results obtained by the two servers. Thus, we did not systematically acquire results from different locations and servers, as this appeared to have no effect.

To answer RQ2, we pooled the top ten URLs from every sampled date for the WEB2013-2014 query set, and we assessed their relevance[4]. Relevance assessments were collected using crowdsourcing. We setup tasks on Amazon Mechanical Turk, assigning each query-website pair to 5 workers. Workers, selected among those with a 90% acceptance rate and at least 1000 tasks completed, were presented with the TREC topic title and description fields and a link to the webpage to be assessed. We simplified the TREC 2013 six-point judgment scale [6] into the following four-point scale: Highly relevant (this point included Nav, Key, and HRel as defined in TREC 2013), Relevant, Not Relevant, and Junk. As suggested in [7], assessments that took less than 4 seconds were discarded. Collected assessments were aggregated as the median of the collected labels for each query-document pair. We analysed the results using ERR@10, nDCG@10, and P@10, as used in the TREC 2013.

3 EMPIRICAL EVALUATION

3.1 Volatility of Search Results

Figure 2 shows the percentage of new URLS introduced on average in the top 10 results. Each square in the heatmap corresponds to a pair of dates, and thus the percentage difference between the results obtained in the two dates. The darker the red tone, the higher the percentage of new URLs being returned on the later date: the diagonal is yellow, indicating no difference (as expected, as we are comparing a date with itself), and we removed the lower part of the heatmap for clarity (the heatmap is symmetric).

Results highlighted in blue refer to the percentage of new URLs retrieved compared to the initial date (the start of our data sampling): in the figure we further visualised this trend as a line plot to give the reader a different representation of the trend and aid interpretation. On average, we found that each day had 24.43% new URLs, compared to the initial sampling date.

[3] azure.microsoft.com/en-au/services/cognitive-services/bing-web-search-api/

[4] We assumed that the content of the linked webpages remained the same throughout our experiments. While this may have not been the case for all results, as reported in [3], only ≈ 1% of urls found in the first two weeks had modified content in the third week of their study.

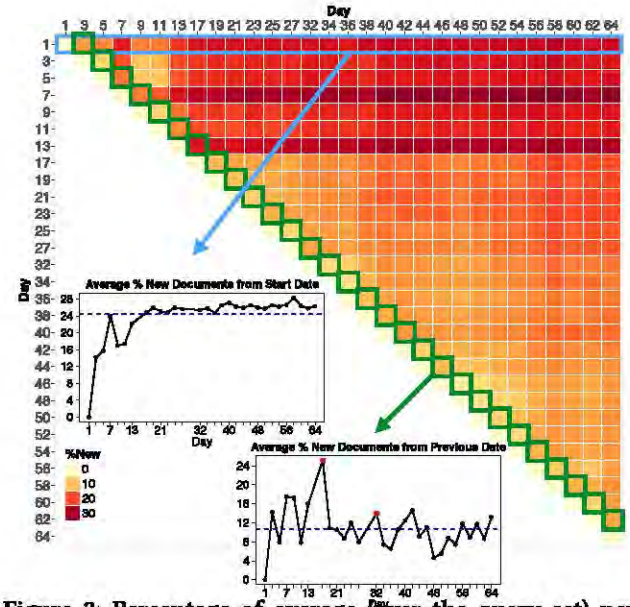

Figure 2: Percentage of average (over the query set) new URLs retrieved in the top 10 results for each pair of sampling dates for the WEB2013-2014 query set.

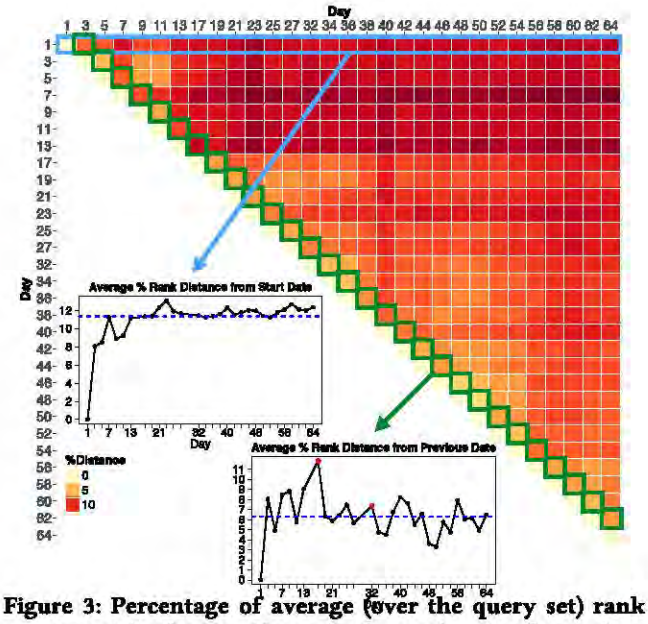

Figure 3: Percentage of average (over the query set) rank movement in the top 10 results for each pair of sampling dates for the WEB2013-2014 query set.

Another important observation from Figure 2 is highlighted in green. This shows the percentage of new urls found on a sampling date compared to the previous sampling date. We further visualised this trend in the corresponding line plot: new urls were found at a comparable percentage overtime. On average, every two days, 10.72% of the URLs retrieved differed from those retrieved on the previous sampling date.[5].

We also investigated the volatility of the top 1 to 9 documents and the top 50 documents, for completeness. We found that percentage volatility trends were similar across rank thresholds (also for CLEF 2016 – results available at http://ielab.io/se-volatility).

We then further investigated the ranking distance between occurrences of the same URL across different dates. Results were again represented as a heatmap, which is reported in Figure 3. The heatmap shows the percentage of rank distances between the top 10 URLs for each pair of sampling dates using the WEB2013-2014 query set. Blue and green colours were used to represent similar circumstances as for the previous heatmap. When considering rank movements over time with respect to the first sampling day (blue highlighting), we found that on average URLs moved 11.36% up or down the ranking, with lesser movement found in the first few days of the experiment. When considering rank movements over time with respect to the previous sampling day (green highlighting), we found that on average URLs moved by 6.29% up or down the ranking compared to the previous date, though peaks with larger rank movements did occur. Interestingly, the line plots in Figures 2 and 3 suggest that the trends observed for new URLs were similar to those for rank distance over time.

Given these results, we answer RQ1 by reporting that, on average, between two consecutive days, search engine results change by 10.72% in terms of new URLs retrieved (1.07 new URLs every 10). Furthermore, we found that the difference is even larger if a wider timespan is considered. In addition, we also report that URLs that occur in the results between two dates are likely to exhibit a rank movement of on average by 6.29%.

[5]Note we missed sampling on the 13, 27, and 29 December 2017: for the sampling date after the ones we missed, rankings were compared to those in the previous available date. In the line plot, we distinguished the data for these dates using red dots.

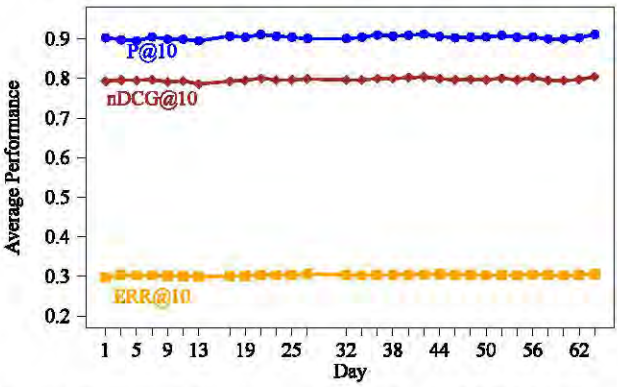

Figure 4: Average effectiveness of the top 10 URLs retrieved for the WEB2013-2014 query set.

3.2 Impact of Result Volatility on Search Effectiveness

Figure 4 shows the search effectiveness over time for the WEB2013-2014 query set, averaged over all queries for each sampling date. The average trends show that search engine volatility had little impact on the average search effectiveness: despite new URLs were retrieved over time, and existing URLs changed rank, effectiveness on average did not vary significantly. Statistical significant differences (t-test $p < 0.05$) were found only between $\approx 6.7\%$ of the results for each pair of days (only unique pairs were considered).

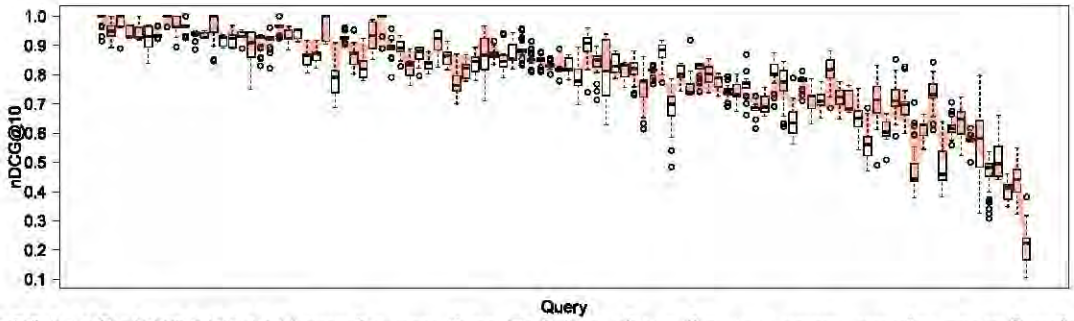

Figure 5: nDCG@10 for WEB2013-2014 queries over time. Each box plot refers to a query; queries are ordered in decreasing effectiveness as returned on day 1. The red shaded area indicates the effectiveness gap between results for day 1 and those for the day with the biggest average effectiveness gap over the query set.

The previous results analysed the impact of search engine volatility by averaging effectiveness over the query set. We next analyse the impact volatility had on a query-by-query basis; we did this for graded relevance (nDCG@10) – similar findings were observed for other settings. Results are reported in Figure 5. Box plots were organised such that queries were ordered in decreasing effectiveness of the results obtained on the initial date of sampling (day 1): each box summarises the effectiveness of a query over time. The box plots show effectiveness did vary over time for each query, with some queries achieving substantially different effectiveness depending on the date. Specifically, we found that 67 out of 100 queries had a change in nDCG@10 that was higher or equal to 0.1 and, on average, individual query effectiveness varied by 0.1431 over the sampling period, with the largest variation recorded being 0.4700. To further provide an intuition of the gap in effectiveness that search result volatility generated, we highlighted the gap between the effectiveness of each query recorded on the first day of our experiment, and the day with the biggest average effectiveness gap over in the query set (day 64). This gap is represented by the red shaded area in Figure 5.

4 DISCUSSION

The findings from our experiments quantified the amount of volatility measured in 2-day time intervals and over longer periods. In addition, they also highlighted that not only are the search engines volatile, but their effectiveness is also volatile, given a query; although we found that for the query set used, average effectiveness remained mostly unchanged.

These findings suggest that search engine result volatility is likely to largely impact the *replicability* of results obtained by exploiting commercial search engine APIs either for algorithmical advances or within user studies. We also argue that volatility also impacts *reproducibility*, as the deterioration of results over time for some queries is large and, if results are used algorithmically, is likely to produce different outcomes and thus affecting the actual quality of techniques like query expansion based on the initial search engine results. While not done here, we aim to empirically investigate this in future work.

While our findings suggest that search engine result volatility may affect results of information retrieval studies, it is unclear how researchers could mitigate these issues and yet use commercial search engines and associated APIs within their research. A possible avenue may be repeating the studies over a sufficiently long period of time, so as to account for search engine volatility as one of the factors affecting results and study this with respect to their results.

5 CONCLUSION

In this paper, we investigated the volatility of commercial search engines and its impact on information retrieval research. By sampling the results returned by the Bing Web Search API every two days for a period of 64 days, we found that, on average, the search engine retrieved 10.72% new URLs in the top 10 ranks. Additionally, we also found that a URL that was retrieved on a previous date was subject to an average rank movement of 6.29%. When examining the possible impact such a volatility may have on information retrieval research that makes use of such search services, we found that on average, nDCG@10 varied by 0.1431 (19.88%) for each query and the biggest nDCG@10 variation for a query was 0.4700 (143.51%). These results suggest that research that uses commercial search engines as part of an algorithmical pipeline or user study should be aware of search engine volatility and its implications.

Acknowledgements. Jimmy is sponsored by the Indonesia Endowment Fund for Education (Lembaga Pengelola Dana Pendidikan / LPDP). Guido Zuccon is the recipient of an Australian Research Council DECRA Research Fellowship (DE180101579). This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 732328.

REFERENCES

[1] Eytan Adar, Jaime Teevan, Susan T Dumais, and Jonathan L Elsas. 2009. The Web Changes Everything: Understanding the dynamics of web content. In *WSDM'09*.
[2] I Altingovde, R Ozcan, and O Ulusoy. 2011. Evolution of Web Search Results within Years. In *SIGIR'11*.
[3] X Bai and F Junqueira. 2012. Online Result Cache Invalidation for Real-Time Web Search. In *SIGIR'12*.
[4] R Blanco, E Bortnikov, F Junqueira, R Lempel, L Telloli, and H Zaragoza. 2010. Caching Search Engine Results over Incremental Indices. In *SIGIR'10*.
[5] R Cilibrasi and P Vitanyi. 2007. The Google Similarity Distance. *TKDE* 19, 3 (2007).
[6] K Collins-Thompson, P Bennett, F Diaz, C Clarke, and E Voorhees. 2014. TREC 2013 Web Track Overview. In *TREC*.
[7] E. Maddalena, M. Basaldella, D. De Nart, D. Degl'Innocenti, S. Mizzaro, and G. Demartini. 2016. Crowdsourcing Relevance Assessments: The unexpected benefits of limiting the time to judge. In *HCOMP'16*.
[8] D Maxwell, L Azzopardi, and Y Moshfeghi. 2017. A Study of Snippet Length and Informativeness: Behaviour, Performance and User Experience. In *SIGIR'17*.
[9] Frank McCown and Michael L Nelson. 2007. Search Engines and their Public Interfaces: Which APIs are the most synchronized?. In *WWW'07*.
[10] M Symonds, P Bruza, G Zuccon, B Koopman, L Sitbon, and I Turner. 2014. Automatic Query Expansion: A structural linguistic perspective. *JAIST* 65, 8 (2014).

ANN ARBOR
MICHIGAN · USA

08-12 JULY
2018

SIGIR #41

PROCEEDINGS

www.sigir.org/sigir2018

Association for
Computing Machinery

The Association for Computing Machinery
2 Penn Plaza, Suite 701
New York, New York 10121-0701

ISBN: 978-1-4503-5657-2

Additional copies may be ordered prepaid from:

ACM Order Department
PO Box 30777
New York, NY 10087-0777, USA

Phone: 1-800-342-6626 (USA and Canada)
+1-212-626-0500 (Global)
Fax: +1-212-944-1318
E-mail: acmhelp@acm.org
Hours of Operation: 8:30 am – 4:30 pm ET

Printed in the USA

ACM SIGIR 2018 Chairs' Welcome

Welcome to the 41st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2018), the premier scientific conference in the broad area of information retrieval. We're delighted to host ACM SIGIR 2018 on the beautiful University of Michigan campus in Ann Arbor, U.S.A. We received a high-quality set of submissions for full and short papers, demos, industry papers, tutorials, workshops, and doctoral consortium, to consider for inclusion in the overall conference program. In particular we thank the 15 Track co-Chairs, 93 Senior Program Committee (SPC) members, 343 Program Committee (PC) members, and 91 additional reviewers for their contributions to full and short paper selection. This pool of committed SIGIR volunteers was based in 35 countries/regions and over 255 institutions, spanning academia, industry, and beyond.

We reviewed 409 papers for the full paper track, and accepted 86, for an acceptance rate of 21%. In total, there were 1247 authors from 33 countries: 440 from China, 325 from the United States, and 482 from the rest of the world. The top five countries in terms of accepted papers (taking all author affiliations of each paper equally into account) were China (34%), the U.S.A. (30%), Australia (6%), Canada (4%) and Germany (4%). The top three countries in terms of acceptance rates (taking all author affiliations of each paper equally into account) were Australia (44%), Israel (32%) and Germany (31%). The most popular track for submissions was "Content Recommendation, Analysis and Classification" (124 submissions) while the tracks with highest representation in the accepted papers were "Human Factors and Interfaces" and "Domain-Specific Applications" (both 20 papers).

This is the first time in recent years that we have adopted a track-based structure in paper submission and review processes. This change was intended to help us attract more submissions from closely related research communities, and the introduction of "track chairs" (listed below) also helped recruit more domain experts for the program committee and make high quality judgments for the submissions.

Search and Ranking
Iadh Ounis, University of Glasgow
James Allan, University of Massachusetts Amherst
Ji-Rong Wen, Renmin University of China

Foundations and Future Directions
Fernando Diaz, Spotify
Maarten de Rijke, University of Amsterdam

Domain-Specific Applications
Carlos Castillo, Universitat Pompeu Fabra
Tat-Seng Chua, National University of Singapore
Debora Donato, Mix

Content Recommendation, Analysis and Classification
Yi Zhang, University of California Santa Cruz
Min-Yen Kan, National University of Singapore
Deepak Agarwal, LinkedIn

Artificial Intelligence, Semantics, and Dialog
Jianfeng Gao, Microsoft
Jun Wang, University College London

Human Factors and Interfaces
Jeff Huang, Brown University
Rob Capra, University of North Carolina

For all papers, we first assigned 3 PC members and 1 SPC member to provide the first-round review opinions. After that, if the SPC member believed that the review process would benefit from an additional review, we assigned an additional SPC to provide a second-round opinion. The SPC would then make a preliminary recommendation, and the track chairs reviewed the recommendations in their track. The final acceptance decisions were made at the in-person PC meeting held in Ann Arbor, U.S.A. on April 6-7, 2018. Video conferencing was provided at the PC meeting to allow remote participation of additional SPC members. For the 409 valid submissions, the program committee carefully reviewed all of them and provided at least three reviews for each. An additional 65 reviews were also performed by the SPC members, which gave the corresponding papers comments from at least four reviewers.

We received 327 short paper submissions and accepted 98 (30%) of them. The short paper chairs were Paul Bennett and Min Zhang.

We received 36 submissions for demonstrations and accepted 18 (50%) of them. The demo chairs were Grace Hui Yang and Yue Wang.

This year the program for SIRIP Industry Days, chaired by Yi Chang and Dawei Yin, was increased from one day to two days. Across both days, the program had four keynotes, three invited talks, and 9 (56%) accepted talks from 16 submissions.

The doctoral consortium received 16 submissions of which 12 (75%) were accepted. The DC was chaired by Makoto P. Kato and Charlie Clarke.

Eleven out of thirteen (85%) tutorial proposals were accepted. The tutorial chairs were Mark Smucker and Jian-Yun Nie.

Fifteen workshop proposals were received, of which 10 (67%) were accepted. Hui Fang and Fiana Raiber chaired the workshop review process.

We sincerely thank the above track and committee chairs, their program committee members and reviewers for all their hard work—including the best paper selection committee led by Hang Li. We believe the overall result to be a technical program of outstanding quality that also reflects the wonderful diversity of compelling problems being addressed across the broad field of information retrieval.

Brian D. Davison
SIGIR'18 PC Co-chair
Lehigh University,
United States of America

Emine Yilmaz
SIGIR'18 PC Co-chair
University College London,
United Kingdom

Yiqun Liu
SIGIR'18 PC Co-chair
Tsinghua University,
China

Kevyn Collins-Thompson
SIGIR'18 General Co-chair
University of Michigan,
United States of America

Qiaozhu Mei
SIGIR'18 General Co-chair
University of Michigan,
United States of America

Table of Contents

Session 1D: Learning to Rank I

Session Chair: Hang Li *(Bytedance Technologies)*

Session 2A: Sentiment & Opinion

Session Chair: Min Zhang *(Tsinghua University)*

Session 2B: Social

Session Chair: Vanessa Murdock *(Amazon)*

Session 2C: App Search & Recommendation

Session Chair: Tat-Seng Chua *(National University of Singapore)*

Session 2D: Conversational Systems

Session Chair: Jian-Yun Nie *(University of Montreal)*

Session 3A: Social Good

Session Chair: David Carmel *(Amazon Research)*

Session 3B: Privacy

Session Chair: Grace Hui Yang *(Georgetown University)*

Session 3C: Question Answering

Session Chair: Jianfeng Gao *(Microsoft Research)*

Session 3D: Learning to Rank II

Session Chair: Ji-Rong Wen *(Renmin University of China)*

Session 4A: Fairness & Robustness

Session Chair: Fernando Diaz *(Spotify)*

Session 4B: Behavior

Session Chair: Diane Kelly *(University of Tennessee)*

Session 4C: Medical & Legal IR

Session Chair: Isabelle Moulinier *(Capital One)*

Session 4D: Recommender Systems - Methods

Session Chair: Jun Xu *(Chinese Academy of Sciences)*

Session 5A: Location & Trajectory

Session Chair: Dawei Yin *(JD.com)*

Session 5B: Entities

Session Chair: Arjen de Vries *(Radboud University)*

Session 5C: New Metrics

Session Chair: Tetsuya Sakai *(Waseda University)*

Session 5D: Recommender Systems - Applications

Session Chair: Yi Zhang *(University of California, Santa Cruz)*

Session 6A: Evaluation

Session Chair: Mark Sanderson *(RMIT University)*

Session 6B: Hashing & Embedding

Session Chair: Nick Craswell *(Microsoft)*

Session 6C: Knowledge Bases/Graphs

Session Chair: Krisztian Balog *(University of Stavanger)*

Session 6D: Mobile User Behavior

Session Chair: Makoto P. Kato *(Kyoto University)*

Session 7A: Crowdsourcing & Assessment

Session Chair: Mark Smucker *(University of Waterloo)*

Session 7B: Content & Semantics

Session Chair: James Allan *(University of Massachusetts – Amherst)*

Session 7C: Interfaces

Session Chair: Rob Capra *(The University of North Carolina at Chapel Hill)*

Short Research Papers I

Short Research Papers II

Demonstration Papers I

Demonstration Papers II

SIRIP: Industry Days

Tutorials

Workshops

Doctoral Consortium

ACM SIGIR 2018 Conference Organization

General Chairs	Kevyn Collins-Thompson (University of Michigan, United States) Qiaozhu Mei (University of Michigan, United States)
Program Chairs	Brian Davison (Lehigh University, United States) Yiqun Liu (Tsinghua University, China) Emine Yilmaz (University College London, United Kingdom)
Treasurer	VG Vinod Vydiswaran (University of Michigan, United States)
Short Paper Chairs	Paul Bennett (Microsoft, United States) Min Zhang (Tsinghua University, China)
Track Chairs (Core IR)	James Allan (University of Massachusetts Amherst, United States) Iadh Ounis (University of Glasgow, United Kingdom) Ji-Rong Wen (Renmin University of China, China)
Track Chairs (Foundations)	Fernando Diaz (Spotify, United States) Maarten de Rijke (University of Amsterdam, Netherlands)
Track Chairs (Apps)	Carlos Castillo (Universitat Pompeu Fabra, Spain) Tat-Seng Chua (National University of Singapore, Singapore) Debora Donato (Mix, United States)
Track Chairs (Recommender)	Deepak Agarwal (LinkedIn, United States) Min-Yen Kan (National University of Singapore, Singapore) Yi Zhang (University of California Santa Cruz, United States)
Track Chairs (AI)	Jianfeng Gao (Microsoft, United States) Jun Wang (University College London, United Kingdom)
Track Chairs (HCI)	Rob Capra (University of North Carolina, United States) Jeff Huang (Brown University, United States)
Webmasters	Ryan Burton (University of Michigan, United States) Cristina Garbacea (University of Michigan, United States)
Publicity and Media Chairs	Yubin Kim (Carnegie Mellon University, United States) Ying Ding (Indiana University, United States)

Sponsorship Chair	Jimmy Huang (York University, Canada)
Tutorial Chairs	Jian-Yun Nie (Université de Montréal, Canada) Mark Smucker (University of Waterloo, Canada)
Workshop Chairs	Hui Fang (University of Delaware, United States) Fiana Raiber (Yahoo Research, Israel)
Demo Chairs	Grace Hui Yang (Georgetown University, United States) Yue Wang (University of Michigan, United States)
DC Chairs	Charlie Clarke (Facebook, United States) Makoto P. Kato (Kyoto University, Japan)
Industry Day Chairs	Yi Chang (Huawei, China) Dawei Yin (JD.com, United States)
Local Arrangement Chairs	Alex Kotov (Wayne State University, United States) SungJin Nam (University of Michigan, United States) Todd Stuart (University of Michigan, United States)
Registration Chair	Rohail Syed (University of Michigan, United States)
Student Affairs Chair	Laura Dietz (University of New Hampshire, United States)
Student Sponsorship Chair	Hongning Wang (University of Virginia, United States)
Student Volunteer Chair	Wei Ai (University of Michigan, United States)
Awards Chair	Hang Li (Toutiao, China)
Proceedings Chair	Carsten Eickhoff (Brown University, United States)
Brand and Creative Chair	Esther Eickhoff (Art Basel, United States)

Senior PC Members

Eytan Adar (University of Michigan, United States)
Eugene Agichtein (Emory University, United States)
Amr Ahmed (Google, United States)
Omar Alonso (Microsoft, United States)
Xavier Amatriain (Curai, United States)
Javed Aslam (Northeastern University, United States)
Leif Azzopardi (University of Strathclyde, United Kingdom)
Ricardo Baeza-Yates (NTENT Northeastern University, USA)
Krisztian Balog (University of Stavanger, Norway)
Nicholas Belkin (Rutgers University, United States)
Michael Bendersky (Google, United States)
Robin Burke (Depaul University, United States)
Jamie Callan (Carnegie Mellon University, United States)
David Carmel (Amazon, Israel)
Ben Carterette (University of Delaware, United States)
Yi Chang (Huawei Research America, United States)
Li Chen (Hong Kong Baptist University, Hong Kong)
Nick Craswell (Microsoft, United States)
Bruce Croft (University of Massachusetts Amherst, United States)
Arjen P. de Vries (Radboud University, Netherlands)
Zhicheng Dou (Renmin University of China, China)
David Elsweiler (University of Regensburg, Germany)
Hui Fang (University of Delaware, United States)
Michel Galley (Microsoft, United States)
Julio Gonzalo (UNED, Spain)
Ido Guy (eBay Research, Israel)
Claudia Hauff (Delft University of Technology, Netherlands)
Jiyin He (CWI, Netherlands)
Qi He (LinkedIn, United States)
Djoerd Hiemstra (University of Twente, Netherlands)
Wynne Hsu (National University of Singapore, Singapore)
Xuanjing Huang (Fudan University, China)
Grace Hui Yang (Georgetown University, United States)
Hideo Joho (University of Tsukuba, Japan)
Joemon Jose (University of Glasgow, United Kingdom)
Vanja Josifovski (Pinterest, United States)
Jaap Kamps (University of Amsterdam, Netherlands)
George Karypis (University of Minnesota, United States)
Makoto P. Kato (Kyoto University, Japan)
Diane Kelly (University of Tennessee, United States)
Yehuda Koren (Google, Israel)
Oren Kurland (Technion, Israel Institute of Technology, Israel)
Wai Lam (The Chinese University of Hong Kong, Hong Kong)
Wang-Chien Lee (The Pennsylvania State University, United States)
Ee-Peng Lim (Singapore Management University, Singapore)

Christina Lioma (University of Copenhagen, Denmark)
Tie-Yan Liu (Microsoft, China)
Yoelle Maarek (Amazon, Israel)
Craig Macdonald (University of Glasgow, United Kingdom)
R. Manmatha (A9/Amazon, United States)
Tao Mei (JD AI Research, China)
Donald Metzler (Google, United States)
Prasenjit Mitra (The Pennsylvania State University, United States)
Marie-Francine Moens (Katholieke Universiteit Leuven, Belgium)
Alistair Moffat (The University of Melbourne, Australia)
Isabelle Moulinier (Capital One, United States)
Vanessa Murdock (Microsoft, United States)
Wolfgang Nejdl (L3S and University of Hannover, Germany)
Jian-Yun Nie (University of Montreal, Canada)
Nuria Oliver (Vodafone Research, Spain)
Salvatore Orlando (Università Ca' Foscari Venezia, Italy)
Raffaele Perego (ISTI-CNR, Italy)
Benjamin Piwowarski (CNRS / Sorbonne Université, LIP6, France)
Barbara Poblete (University of Chile, Chile)
Tao Qin (Microsoft, China)
Filip Radlinski (Google, United Kingdom)
Steffen Rendle (Google, United States)
Berthier Ribeiro-Neto (UFMG, Belo Horizonte, Brazil)
Paolo Rosso (Universitat Politècnica de València, Spain)
Tetsuya Sakai (Waseda University, Japan)
Mark Sanderson (RMIT University, Australia)
Falk Scholer (RMIT University, Australia)
Chirag Shah (Rutgers University, United States)
James G. Shanahan (Independent Consultant (San Francisco), United States)
Jialie Shen (Newcastle University, United Kingdom)
Milad Shokouhi (Microsoft, United States)
Fabrizio Silvestri (Facebook, United Kingdom)
Mark Smucker (University of Waterloo, Canada)
Aya Soffer (IBM, Israel)
Ruihua Song (Microsoft, China)
Neel Sundaresan (Microsoft, United States)
Paul Thomas (Microsoft, Australia)
Fei Wang (Cornell University, United States)
Gerhard Weikum (Max Planck Institute for Informatics, Germany)
Ryen W. White (Microsoft, United States)
Lexing Xie (Australian National University, Australia)
Jun Xu (Inst. of Computing Tech., Chinese Academy of Sciences, China)
Yiming Yang (Carnegie Mellon University, United States)
Dawei Yin (JD.com, United States)
Elad Yom-Tov (Microsoft, Israel)

Markus Zanker (Free University of Bozen-Bolzano, Italy)
Dell Zhang (Birkbeck, University of London, United Kingdom)
Weinan Zhang (Shanghai Jiao Tong University, China)

PC Members

Panagiotis Adamopoulos (Emory University, United States)
Qingyao Ai (University of Massachusetts Amherst, United States)
Luca Maria Aiello (Nokia Bell Labs, United Kingdom)
Elif Aktolga (Apple Inc., United States)
Dyaa Albakour (Signal Media, United Kingdom)
Mohammad Aliannejadi (Università della Svizzera italiana, Switzerland)
Ismail Sengor Altingovde (Middle East Technical University, Turkey)
Giambattista Amati (Fondazione Ugo Bordoni, Italy)
Enrique Amigó (UNED, Spain)
Massih-Reza Amini (LIG, France)
Hadi Amiri (Harvard University, United States)
Avi Arampatzis (Democritus University of Thrace, Greece)
Ioannis Arapakis (Telefonica Research, Spain)
Jaime Arguello (The University of North Carolina at Chapel Hill, United States)
Ahmed Hassan Awadallah (Microsoft, United States)
Seyed Ali Bahrainian (Università della Svizzera italiana, Switzerland)
Xiao Bai (Yahoo Research, United States)
Jie Bao (Urban Computing Lab, JD Finance, China)
Zhifeng Bao (RMIT University, Australia)
Alberto Barrón-Cedeño (Qatar Computing Research Institute, Qatar)
Srikanta Bedathur (IBM, India)
Joeran Beel (Trinity College Dublin, Ireland)
Alejandro Bellogín (Universidad Autónoma de Madrid, Spain)
Klaus Berberich (Max Planck Institute for Informatics, Germany)
Jiang Bian (Microsoft, China)
Bodo Billerbeck (Microsoft, Australia)
Lidong Bing (Tencent AI Lab, Shenzhen, China)
Roi Blanco (Amazon Research, Spain)
Paolo Boldi (University of Milan, Italy)
Ilaria Bordino (UniCredit RD, Italy)
Marc Bron (Yahoo Labs, United Kingdom)
Fidel Cacheda (Universidad de A Coruña, Spain)
Deng Cai (Zhejiang University, China)
Fei Cai (National University of Defense Technology, China)
Fazli Can (Bilkent University, Turkey)
Iván Cantador (Universidad Autónoma de Madrid, Spain)
Marc Cartright (Google, United States)
Pablo Castells (Universidad Autónoma de Madrid, Spain)

James Caverlee (Texas AM University, United States)
Praveen Chandar (Spotify Research, United States)
Kevin Chang (University of Illinois at Urbana-Champaign, United States)
Tao Chen (Google, United States)
Hsin-Hsi Chen (National Taiwan University, Taiwan)
Ruey-Cheng Chen (RMIT University, Australia)
Zheng Chen (Microsoft, China)
Zhumin Chen (shandong university, China)
Max Chevalier (IRIT, Université Paul Sabatier, France)
Jean-Pierre Chevallet (University of Grenoble Alpes, France)
Adrian-Gabriel Chifu (Université d'Aix-Marseille, Université de Toulon, CNRS, LIS, Marseille, France.)
Evangelia Christakopoulou (University of Minnesota, United States)
Paul Clough (The University of Sheffield, United Kingdom)
Arman Cohan (Georgetown University, United States)
Michael Cole (LexisNexis, United States)
Gordon Cormack (University of Waterloo, Canada)
Matt Crane (University of Waterloo, Canada)
Paolo Cremonesi (Politecnico di Milano, Italy)
Peng Cui (Tsinghua University, China)
Shane Culpepper (RMIT University, Australia)
Tonya Custis (Thomson Reuters, United States)
Jeffrey Dalton (University of Glasgow, United Kingdom)
Duc-Tien Dang-Nguyen (Dublin City University, Ireland)
Mostafa Dehghani (University of Amsterdam, Netherlands)
Gianluca Demartini (The University of Queensland, Australia)
Hongbo Deng (Google, United States)
Bhuwan Dhingra (Carnegie Mellon University, United States)
Laura Dietz (University of New Hampshire, United States)
B. Taner Dincer (University of Mugla, Turkey)
Doug Downey (Northwestern University, United States)
Kevin Duh (Johns Hopkins University, United States)
Carsten Eickhoff (Brown University, United States)
Tamer Elsayed (Qatar University, Qatar)
Yi Fang (Santa Clara University, United States)
Fuli Feng (National University of Singapore, Singapore)
Ignacio Fernández-Tobías (NTENT, Spain)
Nicola Ferro (University of Padova, Italy)
Marcus Fontoura (Microsoft, United States)
Luca Foschini (Evidation Health, United States)
Adam Fourney (Microsoft, United States)
Edward Fox (Virginia Tech, United States)
Ana Freire (Universitat Pompeu Fabra, Spain)
Ingo Frommholz (University of Bedfordshire, United Kingdom)
Norbert Fuhr (University of Duisburg-Essen, Germany)

Patrick Gallinari (LIP6 - University of Paris 6, France)
Debasis Ganguly (IBM Research, Ireland)
Bin Gao (Microsoft, United States)
Kiran Garimella (Aalto University, Finland)
Shlomo Geva (Queensland University of Technology, Australia)
Anastasia Giachanou (University of Lugano, Switzerland)
Nadav Golbandi (Google, United States)
Marcos Goncalves (Federal University of Minas Gerais, Brazil)
Jesse Gozali (Mobilewalla, Singapore)
Eduardo Graells-Garrido (Telefónica I+D, Chile)
David Graus (FD Mediagroep, Netherlands)
Gregory Grefenstette (IHMC and Biggerpan, Inc, France)
Krishna Gummadi (MPI-SWS, Germany)
Guibing Guo (Northeastern University, China)
Jiafeng Guo (Institute of Computing Technology, China)
Qi Guo (Google, United States)
Cathal Gurrin (Dublin City University, Ireland)
Kishaloy Halder (National University of Singapore, Singapore)
Martin Halvey (University of Strathclyde, United Kingdom)
Allan Hanbury (Vienna University of Technology, Austria)
Preben Hansen (Stockholm University, Sweden)
Morgan Harvey (Northumbria University, United Kingdom)
Xiangnan He (National University of Singapore, Singapore)
Ben He (University of Chinese Academy of Sciences, China)
Daqing He (University of Pittsburgh, United States)
Brent Hecht (Northwestern University, United States)
Chathra Hendahewa (TD Bank Group, Canada)
Orland Hoeber (University of Regina, Canada)
Eduard Hoenkamp (Queensland University of Technology, Australia)
Liangjie Hong (Etsy, United States)
Richang Hong (Hefei University of Technology, China)
Frank Hopfgartner (The University of Sheffield, United Kingdom)
Hsun-Ping Hsieh (National Cheng Kung University, Taiwan)
Xia Ben Hu (Texas AM University, United States)
Qinmin Vivian Hu (Ryerson University, Canada)
Jimmy Huang (University of York, Canada)
Zi Huang (The University of Queensland, Australia)
Minlie Huang (Tsinghua University, China)
Kai Hui (SAP SE, Germany)
Samuel Huston (Google, United States)
Georgiana Ifrim (University College Dublin, Ireland)
Diego Fernández Iglesias (University of A Coruña, Spain)
Muhammad Imran (Qatar Computing Research Institute, Qatar)
Amir Ingber (Amazon, Israel)
Peter Ingwersen (University of Copenhagen, Royal School of LIS, Den-

mark)
Radu Tudor Ionescu (University of Bucharest, Faculty of Mathematics and Computer Science, Romania)
Shoaib Jameel (Kent University, United Kingdom)
Jiepu Jiang (College of Information and Computer Sciences, University of Massachusetts Amherst, United States)
Peng Jiang (Alibaba Group, China)
Gareth Jones (Dublin City University, Ireland)
Rosie Jones (Spotify, United States)
Vasileios Kandylas (Microsoft, United States)
Evangelos Kanoulas (University of Amsterdam, Netherlands)
Jussi Karlgren (Gavagai KTH, Sweden)
Jaana Kekäläinen (University of Tampere, Finland)
Liadh Kelly (Maynooth University, Ireland)
Tom Kenter (Google London, United Kingdom)
Madian Khabsa (Apple, United States)
Evgeny Kharitonov (Facebook, France)
Sunyoung Kim (Rutgers University, United States)
Yubin Kim (Carnegie Mellon University, United States)
Julia Kiseleva (Eindhoven University of Technology, Netherlands)
Deguang Kong (Oath, United States)
Weize Kong (Google, United States)
Alexander Kotov (Wayne State University, United States)
Pigi Kouki (University of California Santa Cruz, United States)
Wessel Kraaij (Leiden University, Netherlands)
Udo Kruschwitz (University of Essex, United Kingdom)
Abhimanyu Lad (linkedin, United States)
Olivier Van Laere (Apple Inc., United States)
Yanyan Lan (ICT, China)
Ni Lao (SayMosaic, United States)
Léa Laporte (INSA Lyon - LIRIS, France)
Neal Lathia (Monzo, United Kingdom)
Edith Law (University of Waterloo, Canada)
Matt Lease (The University of Texas at Austin, United States)
Chia-Jung Lee (Microsoft, United States)
Jin Ha Lee (University of Washington, United States)
Sungjin Lee (Microsoft, United States)
Jochen L. Leidner (Thomson Reuters, United Kingdom)
Luis Leiva (Sciling, SL, Spain)
Yong Li (Tsinghua University, China)
Ruoyu Li (University of Texas at Arlington, United States)
Cheng Li (Google, United States)
Chenliang Li (Wuhan University, China)
Jundong Li (Arizona State University, United States)
Liangda Li (Yahoo Lab!, United States)

Xiujun Li (Microsoft, United States)
Panagiotis Liakos (University of Athens, Greece)
Defu Lian (University of Science and Technology of China, China)
Shangsong Liang (University College London, United Kingdom)
Lizi Liao (National University of Singapore, Singapore)
Alexander Libov (Amazon, Israel)
Daniel Liebling (Google, United States)
Jimmy Lin (University of Waterloo, Canada)
Zachary Lipton (University of California San Diego, United States)
Huan Liu (Arizona State University, United States)
Xiaozhong Liu (Indiana University Bloomington, United States)
Chang Liu (Peking University, China)
Jingjing Liu (University of South Carolina, United States)
Kang Liu (National Laboratory of Pattern Recognition,Institute of Automation, Chinese Academy of Sciences, China)
Qiaoling Liu (CareerBuilder LLC, United States)
Wei Liu (Tencent AI Lab, China)
Xiaodong Liu (Microsoft, United States)
Yang Liu (Wilfrid Laurier University, Canada)
Zhiyuan Liu (Tsinghua University, China)
Andreas Lommatzsch (TU Berlin, Germany)
Wei Lu (Singapore University of Technology and Design, Singapore)
Claudio Lucchese (Ca' Foscari University of Venice, Italy)
Cheng Luo (Tsinghua University, China)
Ping Luo (Institute of Computing Technology, CAS; University of Chinese Acadamy of Sciences, China)
Mihai Lupu (Research Studios Austria, Austria)
Qin Lv (University of Colorado Boulder, United States)
Yuanhua Lv (Microsoft, United States)
Hao Ma (Microsoft, United States)
Jun Ma (School of Computer Science and Technology, Shandong University, China)
Weizhi Ma (Tsinghua University, China)
Andrew Macfarlane (City University of London, United Kingdom)
Joao Magalhaes (Universidade NOVA de Lisboa, Portugal)
Eric Malmi (Google, Switzerland)
Tomohiro Manabe (Yahoo Japan Corporation, Japan)
Jiaxin Mao (Tsinghua University, China)
Xianling Mao (Beijing Institute of Technology, China)
Ilya Markov (University of Amsterdam, Netherlands)
Yosi Mass (IBM Research, Haifa, Israel)
Michael Mathioudakis (University of Helsinki, Finland)
Julian Mcauley (University of California San Diego, United States)
Richard Mccreadie (University of Glasgow, United Kingdom)
Rishabh Mehrotra (University College London, United Kingdom)

Edgar Meij (Bloomberg L.P., United Kingdom)
Peter Mika (Schibsted, Spain)
Bhaskar Mitra (Microsoft, United Kingdom)
Boughanem Mohand (IRIT University Paul Sabatier Toulouse, France)
Mitra Mohtarami (MIT Computer Science and Artificial Intelligence Lab, United States)
Manuel Montes-Y-Gómez (Instituto Nacional de Astrofísica, Óptica y Electrónica, Mexico)
Fred Morstatter (USC Information Sciences Institute, United States)
Alessandro Moschitti (Amazon, USA)
Yashar Moshfeghi (University of Strathclyde, United Kingdom)
Seung-Hoon Na (Chonbuk National University, South Korea)
Marc Najork (Google, United States)
Franco Maria Nardini (ISTI-CNR, Italy)
Jun Ping Ng (Amazon, United States)
Chong-Wah Ngo (City University of Hong Kong, Hong Kong)
Liqiang Nie (Shandong University, China)
Alexandros Ntoulas (LinkedIn, United States)
Douglas Oard (University of Maryland, United States)
Neil O'Hare (Yahoo!, United States)
Alexandra Olteanu (IBM, United States)
Eli Osherovich (Amazon, Israel)
Vito Claudio Ostuni (Pandora, United States)
Jiaul Paik (IIT Kharagpur, India)
Haris Papageorgiou (Institute for Language and Speech Processing, Greece)
Alexandra Papoutsaki (Pomona College, United States)
Aasish Pappu (Yahoo Research, United States)
Dae Hoon Park (Yahoo Research, United States)
Laurence Park (Western Sydney University, Australia)
Gabriella Pasi (Università degli Studi di Milano Bicocca, Italy)
Virgil Pavlu (Northeastern Univ, United States)
Jan Pedersen (eBay, United States)
Jian Pei (Simon Fraser University, Canada)
Baolin Peng (The Chinese University of Hong Kong, Hong Kong)
Fabio Pinelli (Vodafone Italia)
Vassilis Plachouras (Facebook, United Kingdom)
Lahari Poddar (National University of Singapore, Singapore)
Martin Potthast (Leipzig University, Germany)
Pernilla Qvarfordt (FX Palo Alto Laboratory, United States)
Fiana Raiber (Yahoo Research, Israel)
Maya Ramanath (IIT Delhi, India)
Edie Rasmussen (The University of British Columbia, Canada)
Xiang Ren (University of Southern California, United States)
Zhaochun Ren (JD.com, United Kingdom)
Haggai Roitman (IBM Research Haifa, Israel)

Tuukka Ruotsalo (University of Helsinki, Finland)
Alan Said (University of Skövde, Sweden)
Rodrygo Santos (Universidade Federal de Minas Gerais, Brazil)
Denis Savenkov (Facebook, United States)
Ralf Schenkel (Trier University, Germany)
Steven Schockaert (Cardiff University, United Kingdom)
Anne Schuth (De Persgroep, Netherlands)
Pavel Serdyukov (Yandex, Russia)
Huawei Shen (Chinese Academy of Sciences, China)
Yelong Shen (Microsoft, United States)
Bichen Shi (University College Dublin, Ireland)
Yue Shi (Facebook, United States)
Gianmaria Silvello (University of Padova, Italy)
Laurianne Sitbon (Queensland University of Technology, Australia)
Catherine Smith (Kent State University, United States)
Luca Soldaini (Georgetown University, United States)
Dawei Song (The Open University, United Kingdom)
Xuemeng Song (Shandong University, China)
Yang Song (Google, United States)
Akshay Soni (Yahoo, United States)
Alessandro Sordoni (Université Pierre et Marie Curie Paris VI, Canada)
Laure Soulier (Sorbonne Universités UPMC-LIP6, France)
Damiano Spina (RMIT University, Australia)
Thomas Steiner (Google, Germany)
Jannik Strötgen (Max Planck Institute for Informatics, Germany)
Torsten Suel (New York University, United States)
Kazunari Sugiyama (National University of Singapore, Singapore)
Aixin Sun (Nanyang Technological University, Singapore)
Jiliang Tang (Michigan State University, United States)
Jinhui Tang (Nanjing University of Science and Technology, China)
Thrivikrama Taula (Ebay Inc, United States)
Gabriele Tolomei (University of Padua, Italy)
Nicola Tonellotto (ISTI-CNR, Italy)
Christoph Trattner (University of Bergen, Norway)
Johanne R. Trippas (RMIT University, Australia)
Andrew Trotman (University of Otago, New Zealand)
Manos Tsagkias (904Labs, Netherlands)
Ming-Feng Tsai (National Chengchi University, Taiwan)
Kazutoshi Umemoto (The University of Tokyo, Japan)
David Vallet (Google, Australia)
Saúl Vargas (ASOS.com, United Kingdom)
Suzan Verberne (LIACS, Leiden University, Netherlands)
Bin Wang (Institute of Information Engineering, Chinese Academy of Sciences, China)
Hongning Wang (University of Virginia, United States)

Jian Wang (Google, United States)
Jingang Wang (Alibaba Group, China)
Lidan Wang (University of Maryland, United States)
Meng Wang (Hefei University of Technology, China)
Quan Wang (Institute of Information Engineering, Chinese Academy of Sciences, China)
Shuaiqiang Wang (JD.COM, China)
Suhang Wang (Arizona State University, United States)
Weiqing Wang (University of Queensland, Australia)
Xuanhui Wang (Google, United States)
Yue Wang (University of Michigan, United States)
Wouter Weerkamp (904Labs, Netherlands)
Jason D. Williams (Microsoft Research, United States)
Kyle Williams (Microsoft, United States)
Liang Wu (Arizona State University)
Chenyan Xiong (Carnegie Mellon University, United States)
Peilin Yang (Twitter Inc., United States)
Ting Yao (Microsoft, China)
Hongzhi Yin (The University of Queensland, Australia)
Zhengtao Yu (Kunming University of Science and Technology, China)
Xiaojun Yuan (University at Albany, State University of New York, United States)
Reza Zafarani (Syracuse University, United States)
Hamed Zamani (University of Massachusetts Amherst, United States)
Meike Zehlike (Technische Universität Berlin, Germany)
Yongfeng Zhang (Rutgers University, United States)
Zhiwei Zhang (Facebook Inc, United States)
Amy Zhang (Massachusetts Institute of Technology, United States)
Aston Zhang (Amazon AI, United States)
Baichuan Zhang (Facebook, United States)
Fuzheng Zhang (Microsoft, China)
Hanwang Zhang (Nanyang Technological University, Singapore)
Peng Zhang (Tianjin University, China)
Qi Zhang (Fudan University, China)
Sicong Zhang (Georgetown University, United States)
Weinan Zhang (Shanghai Jiao Tong University, China)
Yue Zhang (Singapore University of Technology and Design, Singapore)
Le Zhao (Houzz Inc., United States)
Jun Zhao (National Laboratory of Pattern Recognition,Institute of Automation, Chinese Academy of Sciences, China)
Xin Zhao (Renmin University of China, School of Information, China)
Shiwan Zhao (IBM Research, China)
Yong Zheng (Illinois Institute of Technology, United States)
Ke Zhou (University of Nottingham, United Kingdom)
Feida Zhu (Singapore Management University, Singapore)

Yadong Zhu (Mobvista, China)
Nivio Ziviani (Universidade Federal de Minas Gerais, Brazil)
Guido Zuccon (Queensland University of Technology, Australia)

Additional Reviewers

Khaled Ammar (Thomson Reuters, Canada)
Ghazaleh Beigi (Arizona State University, United States)
Miguel Ángel Álvarez Carmona (Instituto Nacional de Astrofísic, Mexico)
Saurabh Chakravarty (Virginia Tech, United States)
Huimin Chen (Tsinghua University, China)
Jia Chen (Carnegie Mellon University, United States)
Tong Chen (The University of Queensland, Australia)
Xiang Chen (National University of Singapore, Singapore)
Xu Chen (Tsinghua University, China)
Lu Cheng (Arizona State University, United States)
Lingyang Chu (Simon Fraser University, Canada)
Ovidiu Dan (Lehigh University, United States)
Yashar Deldjoo (Polytechnic University of Milan, Italy)
Tyler Derr (Michigan State University, United States)
Kaize Ding (Arizona State University, United States)
Dennis Dosso (University of Padua, Italy)
Sebastien Fournier (LSIS, France)
Luke Gallagher (RMIT University, Australia)
Zheng Gao (Indiana University, United States)
Saptarshi Ghosh (Indian Institute of Technology Kharagpur, India)
Pawan Goyal (IIT Kharagpur, India)
Xinyu Guan (Xi'an Jiaotong University, China)
Chun Guo (Pandora, United States)
Ruocheng Guo (Arizona State University, United States)
Ziyi Guo (JD.com, United States)
Christophe Van Gysel (Apple, Netherlands)
Xu Han (Tsinghua University, China)
Maram Hasanain (Qatar University, Qatar)
Heng Huang (University of Texas at Arlington, United States)
Swayambhoo Jain (Technicolor, United States)
Shiyu Ji (University of California, Santa Barbara, United States)
Zhuoren Jiang (Sun Yat-sen University, China)
Hamid Karimi (Michigan State University, United States)
Mucahid Kutlu (Qatar University, Qatar)
Da Li (University of Missouri, United States)
Liuqing Li (Virginia Tech, United States)
Pengfei Li (RMIT University, Australia)
Wei Li (Dublin City University, Ireland)
Yang Li (Northwestern Polytechnical University, China)

Yuanpeng Li (Baidu USA, United States)
Ruiyuan Li (JD.com, United States)
Yankai Lin (Tsinghua University, China)
Haochen Liu (Michigan State University, United States)
Xiaorui Liu (Michigan State University, United States)
Yunzhong Liu (Oath, United States)
Miao Lu (Oath, United States)
Sidi Lu (Shanghai Jiaotong University, China)
Yadan Luo (The University of Queensland, Australia)
Yao Ma (Michigan State University, United States)
Sean Macavaney (Georgetown University, United States)
Maria Maistro (University of Padua, Italy)
Stefano Marchesin (University of Padua, Italy)
Stefania Marrara (C2T, Italy)
Mehrnoosh Mirtaheri (University of Southern California, United States)
Tahora H. Nazer (Arizona State University, United States)
Massimo Nicosia (University of Trento, Italy)
Joao Palotti (Vienna University of Technology, Austria)
Lahari Poddar (National University of Singapore, Singapore)
Zhi Qiao (IBM Research Lab, China)
Vineeth Rakesh (Arizona State University, United States)
Kan Ren (Xi'an Jiaotong University, China)
Andrés Rosso-Mateus (National University of Colombia, Colombia)
Sijie Ruan (JD.com, United States)
Javier Sánchez-Junquera (Instituto Nacional de Astrofísic, Mexico)
Shubhra Kanti Karmaker Santu (University of Illinois at Urbana Champaign, United States)
David Semedo (Universidade NOVA de Lisboa, Portugal)
Kai Shu (Arizona State University, United States)
Aya Soffer (IBM, Israel)
Lukas Tencer (Pandora, United States)
Junfeng Tian (East China Normal University, China)
Dominika Tkaczyk (ADAPT, Ireland)
Cagri Toraman (Bilkent University, Turkey)
Cunchao Tu (Tsinghua University, China)
Antonio Uva (UniTrento, Italy)
Hossein Vahabi (Pandora, United States)
Manisha Verma (University College London, United Kingdom)
Marco Viviani (University of Milano-Bicocca, Italy)
Shaun Wallace (Brown University, United States)
Qinyong Wang (The University of Queensland, Australia)
Weiqing Wang (The University of Queensland, Australia)
Xiang Wang (National University of Singapore, Singapore)
Yingzi Wang (University of Science and Technology of China, China)
Zhiwei Wang (Michigan State University, United States)

Yu Yang (Simon Fraser University, Canada)
Zaihan Yang (Lehigh University, United States)
Jiawen Yao (University of Texas at Arlington, United States)
Zijun Yao (Rutgers University, United States)
Fajie Yuan (University of Glasgow, United Kingdom)
Xuan Zhang (Virginia Tech, United States)
Yazhou Zhang (Tianjin University, China)
Xiangyu Zhao (Michigan State University, United States)

Tutorial PC Members

Jaime Arguello (The University of North Carolina at Chapel Hill, United States)
Ben Carterette (University of Delaware/Spotify, United States)
Gordon Cormack (University of Waterloo, Canada)
Jeffrey Dalton (University of Glasgow, United Kingdom)
Luanne Freund (The University of British Columbia, Canada)
Claudia Hauff (Delft University of Technology, Netherlands)
Evangelos Kanoulas (University of Amsterdam, Netherlands)
Alistair Moffat (The University of Melbourne, Australia)
Jian-Yun Nie (University of Montreal, Canada)
Yongli Ren (RMIT University, Australia)
Laurianne Sitbon (Queensland University of Technology, Australia)
Mark D. Smucker (University of Waterloo, Canada)
Ian Soboroff (NIST, United States)
Torsten Suel (New York University, United States)
Ellen Voorhees (NIST, United States)

Workshop PC Members

Nicola Ferro (University of Padova, Italy)
Ingo Frommholz (University of Bedfordshire, United Kingdom)
Cathal Gurrin (Dublin City University, Ireland)
Frank Hopfgartner (The University of Sheffield, United Kingdom)
Jing Jiang (Singapore Management University, Singapore)
Gareth Jones (Dublin City University, Ireland)
Diane Kelly (University of Tennessee, United States)
Oren Kurland (Technion, Israel Institute of Technology, Israel)
Hang Li (Toutiao AI Lab, China)
Jian-Yun Nie (University of Montreal, Canada)
Dan Pelleg (Yahoo!, Israel)
Haggai Roitman (IBM Research Haifa, Israel)
Andrew Trotman (University of Otago, New Zealand)
Grace Hui Yang (Georgetown University, United States)

Demo PC Members

Krisztian Balog (University of Stavanger, Norway)
Zhuyun Dai (Carnegie Mellon University, United States)
Zhicheng Dou (Renmin University of China, China)
Carsten Eickhoff (Brown University, United States)
Sumio Fujita (Yahoo! JAPAN Research, Japan)
Debasis Ganguly (Dublin City University, Ireland)
Chase Geigle (University of Illinois at Urbana-Champaign, United States)
Grant Ingersoll (LucidWorks, United States)
Sarvnaz Karimi (CSIRO, Australia)
Yubin Kim (Carnegie Mellon University, United States)
Julia Kiseleva (Eindhoven University of Technology, Netherlands)
Udo Kruschwitz (University of Essex, United Kingdom)
Jung-Tae Lee (Naver Corp., South Korea)
Teerapong Leelanupab (King Mongkut's Institute of Technology Ladkrabang, Thailand)
Jochen L. Leidner (Thomson Reuters, United Kingdom)
Shuai Li (University of Cambridge, United Kingdom)
Wei Li (CNGL, Dublin City University, Ireland)
Jiyun Luo (Georgetown University, United States)
Jiaqi Ma (University of Michigan, United States)
Graham Mcdonald (University of Glasgow, United Kingdom)
Rishabh Mehrotra (University College London, United Kingdom)
Javier Parapar (IRLab, Computer Science Dept., University of A Coruña, Spain)
Vassilis Plachouras (Facebook, United Kingdom)
Davood Rafiei (University of Alberta, Canada)
Andrew Salway (Uni Research, Norway)
Anne Schuth (De Persgroep, Belgium)
Kazunari Sugiyama (National University of Singapore, Singapore)
Rohail Syed (University of Michigan, United States)
Yukihiro Tagami (Yahoo Japan Corporation, Japan)
Kazutoshi Umemoto (The University of Tokyo, Japan)
Onur Varol (Northeastern University, United States)
Bin Wang (Institute of Information Engineering, Chinese Academy of Sciences, China)
Sheng Wang (University of Illinois at Urbana-Champaign, United States)
Shiyan Yan (University of Michigan, United States)
Weinan Zhang (Shanghai Jiao Tong University, China)
Yinan Zhang (University of Illinois at Urbana-Champaign, United States)
Guido Zuccon (Queensland University of Technology, Australia)

DC PC Members

Eugene Agichtein (Emory University, United States)
Jaime Arguello (The University of North Carolina at Chapel Hill, United

States)
Krisztian Balog (University of Stavanger, Norway)
Hannah Bast (University of Freiburg, Germany)
Paul Bennett (Microsoft, United States)
David Carmel (Amazon, Israel)
Ben Carterette (University of Delaware/Spotify, United States)
Charles Clarke (Facebook, United States)
Shane Culpepper (RMIT University, Australia)
Brian Davison (Lehigh University, United States)
Maarten de Rijke (University of Amsterdam, The Netherlands)
Luanne Freund (The University of British Columbia, Canada)
Jimmy Huang (University of York, Canada)
Jaap Kamps (University of Amsterdam, The Netherlands)
Noriko Kando (National Institute of Informatics, Japan)
Makoto P. Kato (Kyoto University, Japan)
Diane Kelly (University of Tennessee, United States)
Oren Kurland (Technion, Israel Institute of Technology, Israel)
Douglas Oard (University of Maryland, United States)
Fiana Raiber (Yahoo Research, Israel)
Pavel Serdyukov (Yandex, Russia)
Mark Smucker (University of Waterloo, Canada)
Ian Soboroff (NIST, United States)
Ji-Rong Wen (Renmin University of China, China)

Sponsors & Supporters

Sponsors

Diamond Supporter

Platinum Supporter

Gold Supporters

Silver Supporters

Bronze Supporters

Exhibitors

Travel Grant

Bloomberg®

Student Luncheon Supporter

facebook

Additional Supporters

SCHOOL OF INFORMATION UNIVERSITY OF MICHIGAN

NSF