ECIR 2015: 37th European Conference on Information Retrieval

Allan Hanbury
TU Wien, Austria
allan.hanbury@tuwien.ac.at

Gabriella Kazai
Semion Ltd, UK
gabriella.kazai@gmail.com

Andreas Rauber
TU Wien, Austria
rauber@ifs.tuwien.ac.at

Norbert Fuhr
University of Duisburg-Essen, Germany
norbert.fuhr@uni-due.de

Abstract

The 37th European Conference on Information Retrieval took place from the 29th of March to the 2nd of April 2015 in Vienna, Austria. This report summarises the conference in terms of the keynotes presented, programme, industry day, tutorials, workshops and student support.

1 Introduction

The 37th European Conference on Information Retrieval took place from the 29th of March to the 2nd of April 2015 in Vienna, Austria. The conference was organised by the Institute of Software Technology and Interactive Systems of the TU Wien, supported by the Austrian Computer Society.

The conference took place in the lecture theatres of the TU Wien, in the 4th district of Vienna, close to the centre of the city. The conference dinner was held in the sumptuous rooms of *Palais Niederösterreich* in the centre of Vienna (Figure 1). A total of 230 people attended the ECIR 2015.

2 Keynotes

ECIR 2015 included three keynote talks (Figure 2):

- Marti Hearst (University of California, Berkeley) — *Still Haven't Found What I'm Looking For: Suggestions for Search Research*
- Ryen W. White (Microsoft Research) — *Mining and Modeling Online Health Search*
- Stefan Thurner (Medical University of Vienna) — *What to Do If You Know Everything? Studying Human Behavior in a Virtual World*

Figure 1: ECIR Conference Dinner at Palais Niederösterreich, Vienna

Figure 2: The three keynote speakers: Marti Hearst, Ryen W. White and Stefan Thurner

Marti Hearst presented a well-structured list of search problems that have not been solved yet, the long tail of research challenges. These include book, software, courseware and

scientific search, as well as effective user interfaces to facilitate collaborative search. These problems will likely serve as inspiration for many of the students attending the ECIR.

The second keynote talk was allocated for the recipient of the 2014 Karen Spärck Jones Award. This award was created by the Information Retrieval Specialist Group of the British Computer Society (BCS IRSG) in conjunction with the BCS to commemorate the achievements of Karen Spärck Jones. This annual award is to encourage and promote talented researchers who have endeavoured to advance our understanding of IR and natural language processing with significant experimental contributions. The 2014 KSJ Award was given to Ryen W. White, who presented an overview of his extensive work on the behaviour of users searching for health and medical information on the web.

In Information Retrieval, we usually have to make many assumptions about the context of a search or of the user of a recommender system. The third talk presented the complete opposite to this situation. Stefan Thurner presented the work that he and his group have done on analysing an online multiplayer game in which they know everything about all interactions between all players, including who is friends with whom, who has spoken to whom, but also who is enemies with whom. Using this data, they were able to produce some ground breaking results in computational social science, some of which were presented in the talk.

3 Programme

ECIR 2015 received a total of 305 submissions in three categories: 190 full papers, 103 short papers, and 12 demonstrations. The geographical distribution of the submissions was as follows: 54% were from Europe, 18% from Asia, 17% from North and South America, 5% from Australasia, and 6% from North Africa and the Middle East. All submissions were reviewed by at least three members of an international two-tier Program Committee. Of the full papers submitted to the conference, 44 were accepted for oral presentation (23%). Of the short papers submitted to the conference, 39 were accepted for poster presentation (38%). In addition, seven demonstrations (58%) were accepted. The full conference proceedings are published in Springer LNCS [2].

3.1 Programme Committee

The Programme Committee consisted of 227 members, of which 31 were Meta-Reviewers and 196 were Reviewers. 62 Additional Reviewers entered some reviews on behalf of a Reviewer. The full paper decisions were made based on a two tier review process involving both Reviewers and Meta-Reviewers, while the decisions on the short papers and demos were made based only on reviews.

Overall, 611 reviews were entered into the review system, amounting to 1 422 024 characters in total. An interesting statistic is the total number of characters entered for all reviews by each reviewer, shown in Figure 3, with black points representing reviewers and blue points representing meta-reviewers. The largest number of characters entered by a single reviewer was around 25 000, while the smallest number of characters entered was around 100.

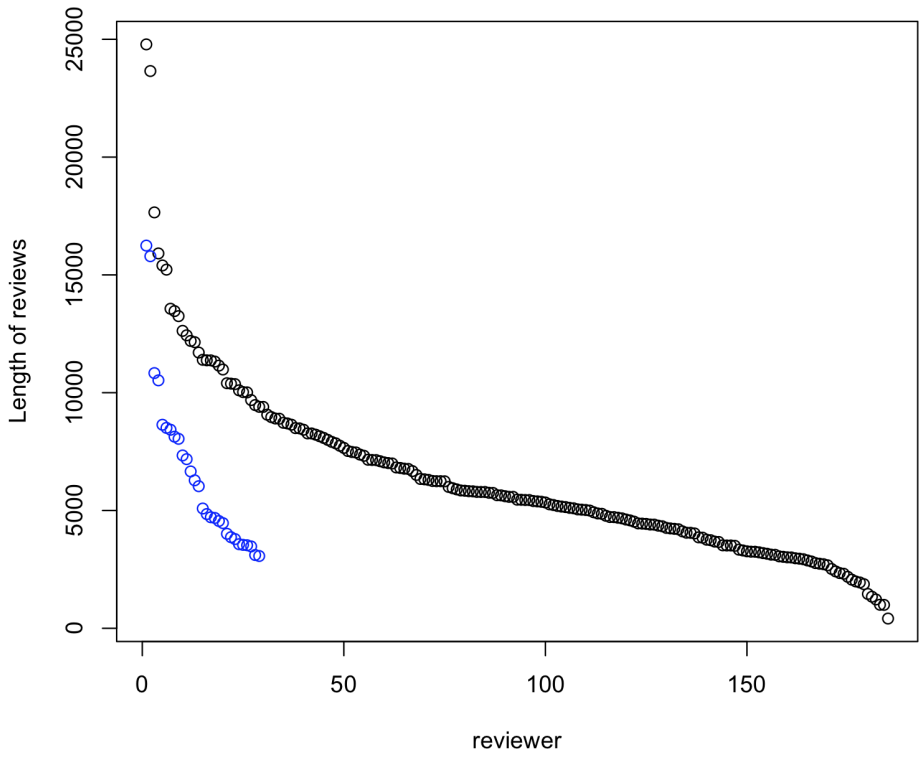

Figure 3: The total number of characters entered for all reviews by each reviewer, in descending order. Black points represent reviewers and blue points represent meta-reviewers.

3.2 Reproducible IR Research Track

An innovation of the ECIR 2015 was the Reproducible IR Research Track. Reproducibility is key for establishing research to be reliable, referenceable and extensible for the future. Experimental papers are therefore most useful when their results can be tested and generalized by peers. This track specifically invited submission of papers reproducing a single or a group of papers in which the authors of the submitted paper had not been involved. It was requested that submitted papers emphasize the motivation for selecting the paper(s), the process of how results have been attempted to be reproduced (successful or not), the communication that was necessary to gather all information, the potential difficulties encountered and the result of the process. A successful reproduction of the work was not a requirement, but it was important to provide a clear and rigorous evaluation of the process to allow lessons to be learned for the future.

Seven papers were submitted to this track, of which three were accepted for oral presentation in a special session on the first day of the conference. After the presentation of the paper, a panel made up of Norbert Fuhr, Jaap Kamps, Diane Kelly and Udo Kruschwitz and chaired by Andreas Rauber debated on *Evaluating IR Research: A Technical or Social Science?*. There was lively participation from the audience in this debate covering the advantages and disadvantages of evaluation approaches ranging from the classic Cranfield paradigm to user studies. The importance of reproducibility of results was a central theme of this debate.

This track was considered to make a strong contribution to the development of an increasingly important topic in Information Retrieval (and computer science in general), and will be continued at the ECIR 2016.

3.3 Awards

A committee led by John Tait (Chair), with members Norbert Fuhr, Donna Harman, Victor Lavrenko, Arjen de Vries and Emine Yilmaz made the decision about the ECIR 2015 *Best Paper Award*. This award was presented to Horatiu Bota, Ke Zhou and Joemon Jose for their paper *Exploring Composite Retrieval from the Users' Perspective*.

The Best Poster and Best Demo awards were chosen by the attendees through a vote. The *Best Poster Award* was presented to Mostafa Dehghani, Hosein Azarbonyad, Maarten Marx and Jaap Kamps for the paper *Sources of Evidence for Automatic Indexing of Political Texts*. The *Best Demo Award* was presented to Miguel Martinez-Alvarez, Udo Kruschwitz, Wesley Hall and Massimo Poesio for their demo *Signal: Advanced Real-Time Information Filtering*.

Awards for the best reviewers and best meta-reviewers were also presented. Awards for the *Best Reviewers* were presented to Jeroen Vuurens (Delft University of Technology) and Matthias Hagen (Bauhaus University Weimar), while awards for the *Best Meta-Reviewers* were presented to Jaime Arguello (University of North Carolina at Chapel Hill) and Stefano Mizzaro (University of Udine).

4 Industry Day

In contrast to industry days in the last few editions of the ECIR, which were held in parallel to the conference sessions, the ECIR 2015 industry day was a stand-alone event held on the day after the main conference programme. The 2015 Industry Day was chaired by Jussi Karlgren and Paul Ogilvie. A further departure from the traditional ECIR industry days was the focus on start-ups and smaller niche companies in the Information Retrieval area, instead of on large established search companies.

The Industry Day began with a keynote talk by William Stevens of Europe Unlimited, a company specialising in bringing together entrepreneurs with venture capital. This was followed by a lively discussion of various aspects of running a start-up, from investment to exit strategies. Representatives of eight companies then gave short presentations, with the age of the companies ranging from recently founded (Lumi, Signal, Unbabel) through established (Precognox, Spinque, Seznam, Catalyst) to long established (Thomson Reuters). The Industry Day way well attended, with over 70 people attending at the peak.

5 Tutorials

The tutorial selection committee was chaired by Birger Larsen and Adrian Iftene. The following five tutorials were selected from the eight proposals submitted. More detailed descriptions of the tutorials are available in the ECIR Proceedings.

5.1 A Formal Approach to Effectiveness Metrics for Information Access: Retrieval, Filtering, and Clustering

This tutorial was presented by Enrique Amigó, Julio Gonzalo and Stefano Mizzaro. They presented, reviewed, and compared the most popular evaluation metrics for some of the most salient information related tasks, covering: (i) Information Retrieval, (ii) Clustering, and (iii) Filtering. The tutorial made a special emphasis on the specification of constraints for suitable metrics in each of the three tasks, and on the systematic comparison of metrics according to how they satisfy such constraints. This comparison provides criteria to select the most adequate metric or set of metrics for each specific information access task. The last part of the tutorial investigated the challenge of combining and weighting metrics.

5.2 Measuring Document Retrievability

This tutorial was presented by Leif Azzopardi. Retrievability is an important and interesting indicator that can be used in a number of ways to analyse Information Retrieval systems and document collections. Rather than focusing totally on relevance, retrievability examines what is retrieved, how often it is retrieved, and whether a user is likely to retrieve it or not. This is important because a document needs to be retrieved, before it can be judged for relevance. In this tutorial, the concept of retrievability along with a number of retrievability measures, how it can be estimated and how it can be used for analysis were explained. Since retrieval precedes relevance, an overview of how retrievability relates to effectiveness – describing some of the insights that researchers have discovered thus far, was also provided. The tutorial also showed how retrievability relates to efficiency, and how the theory of retrievability can be used to improve both effectiveness and efficiency. This was followed by an overview of the different applications of retrievability such as Search Engine Bias, Corpus Profiling, etc., before wrapping up with challenges and opportunities. The final session of the day looked at example problems and ways to analyse and apply retrievability to other problems and domains.

5.3 Visual Analytics for Information Retrieval Evaluation

Nicola Ferro and Giuseppe Santucci presented this tutorial. Measuring is a key to scientific progress. This is particularly true for research concerning complex systems, whether natural or human-built. Multilingual and multimedia information systems are increasingly complex: they need to satisfy diverse user needs and support challenging tasks. Their development calls for proper and new evaluation methodologies to ensure that they meet the expected user requirements and provide the desired effectiveness.

The tutorial introduced basic and intermediate concepts about laboratory-based evaluation of Information Retrieval systems, its pitfalls, and shortcomings and complemented them with a recent and innovative angle to evaluation: the application of methodologies and tools coming from the visual analytics domain for better interacting, understanding, and exploring the experimental results and IR system behaviour.

5.4 Join the Living Lab: Evaluating News Recommendations in Real-time

Frank Hopfgartner and Torben Brodt presented this tutorial. Participants learned how to participate in CLEF-NEWSREEL, a living lab on the evaluation of news recommender algorithms. Various research challenges can be addressed within CLEF-NEWSREEL, including the development and evaluation of collaborative filtering and content based filtering strategies. Satisfying information needs by techniques including preference elicitation, pattern recognition, and prediction, recommender systems connect the research areas Information Retrieval and machine learning.

5.5 Statistical Power Analysis for Sample Size Estimation in Information Retrieval Experiments with Users

This tutorial was presented by Diane Kelly. One critical decision researchers must make when designing laboratory experiments with users is how many participants to study. In interactive Information Retrieval, the determination of sample size is often based on heuristics and limited by practical constraints such as time and finances. As a result, many studies are underpowered and it is common to see researchers make statements like "With more participants significance might have been detected," but what does this mean? What does it mean for a study to be underpowered? How does this affect what we are able to discover, how we interpret study results and how we make choices about what to study next? How does one determine an appropriate sample size? What does it even mean for a sample size to be appropriate? This tutorial addressed these questions by introducing participants to the use of statistical power analysis for sample size estimation in laboratory experiments with users. In discussing this topic, the issues of effect size, Type I and Type II errors and experimental design, including choice of statistical procedures, were addressed.

6 Workshops

The workshop selection committee was chaired by Guido Zuccon and András Benczur. The following five workshops were selected from the seven proposals submitted. More detailed descriptions of the workshops are available in the ECIR Proceedings.

6.1 Second International Workshop on Gamification for Information Retrieval (GamifIR'15)

This workshop was organised by Frank Hopfgartner, Gabriella Kazai, Udo Kruschwitz, Michael Meder and Mark Shovman. Gamification is a popular methodology describing the trend of applying game design principles and elements, such as feedback loops, points, badges or leader boards in non-gaming environments. Gamification can have several different objectives. Besides just increasing the fun factor, these could be, for example, to achieve more accurate work, better retention rates and more cost effective solutions by relating motivations for participating as more intrinsic than conventional methods. In the context of Information Retrieval (IR), there are various tasks that can benefit from gamification techniques. Examples include the manual annotation of documents in IR evaluation and participation in

user studies to tackle interactive IR challenges. Gamification, however, comes with its own challenges and its adoption in IR is still in its infancy.

The Second International Workshop on Gamification for Information Retrieval (GamifIR'15) focused on the challenges and opportunities that gamification can present for the IR community. The workshop brought together researchers and practitioners from a wide range of areas including game design, Information Retrieval, human-computer interaction, computer games, and natural language processing. The proceedings are available in CEUR [3].

6.2 Supporting Complex Search Tasks

This workshop was organised by Maria Gade, Mark Hall, Hugo Huurdeman, Jaap Kamps, Marijn Koolen, Mette Skov, Elaine Toms and David Walsh. There is broad consensus in the field of IR that search is complex in many use cases and applications, both on the Web and in domain specific collections, and both professionally and in daily life. Yet our understanding of complex search tasks, in comparison to simple look up tasks, is fragmented at best.

The workshop addressed the many open research questions: What are the obvious use cases and applications of complex search? What are essential features of work tasks and search tasks to take into account? And how do these evolve over time? With a multitude of information, varying from introductory to specialized, and from authoritative to speculative or opinionated, when to show what sources of information? How does the information seeking process evolve and what are relevant differences between different stages? With complex task and search process management, blending searching, browsing, and recommendations, and supporting exploratory search to sensemaking and analytics, UI and UX design pose an overconstrained challenge. How do we know that our approach is any good?

Supporting complex search task requires new collaborations across the whole field of IR, and the proposed workshop brought together a diverse group of researchers to work together on one of the greatest challenges of the field. The proceedings are available in CEUR [1].

6.3 Fifth Workshop on Context-awareness in Retrieval and Recommendation (CaRR 2015)

This workshop was organised by Ernesto William De Luca, Alan Said, Fabio Crestani and David Elsweiler. Context-aware information is widely available in various ways such as interaction patterns, devices, annotations, query suggestions and user profiles and is becoming more important for enhancing retrieval performance. At the moment, the main issue to cope with is not only retrieving the most relevant items and content, but defining them ad hoc. Further relevant issues are personalizing and adapting the information and the way it is displayed to the user's current situation (device, location, social surrounding) and interests.

The workshop focused on integration notions of social context into retrieval and recommendation. Context is seen as a general factor regarding the user, item, system, etc. e.g. location, weather, mood. The need of personalizing and adapting information is accentuated when we consider this kind of device- and interaction-based context. The CaRR Workshop invited the community to a discussion on new, creative ways to handle context-awareness. Furthermore, CaRR worked toward improving the exchange of ideas between different communities involved in research concerning, among other Information Retrieval, recommender systems, web mining, machine learning, data mining, HCI, etc.

6.4 Second International Workshop on Bibliometric-enhanced Information Retrieval

This workshop was organised by Philipp Mayr, Ingo Frommholz, Peter Mutschke and Andrea Scharnhorst. Bibliometric techniques are not yet widely used to enhance retrieval processes in digital libraries, although they offer value-added effects for users. This workshop explored how statistical modelling of scholarship, such as Bradfordizing or network analysis of coauthorship network, or simple citation graphs, can improve retrieval services for specific communities, as well as for large, cross-domain collections like Mendeley. This workshop raised awareness of the missing link between Information Retrieval (IR) and bibliometrics/scientometrics and created a common ground for the incorporation of bibliometric-enhanced services into retrieval at the scholarly search engine interface. The second BIR workshop (BIR2015) had the goal of applying insights from bibliometrics, scientometrics, and informetrics to concrete, practical problems of Information Retrieval and browsing. The proceedings are available in CEUR [4].

6.5 Multimodal Retrieval in the Medical Domain (MRMD 2015)

This workshop was organised by Henning Müller, Oscar Jiménez Del Toro, Allan Hanbury, Georg Langs and Antonio Foncubierta-Rodríguez. Medical information is of interest to a wide variety of users, including patients and their families, researchers, general practitioners and clinicians, and practitioners with specific expertise such as radiologists. There are several dedicated services that seek to make this information more easily accessible, such as Health on the Net's medical search systems for the general public and medical practitioners (`http://www.hon.ch/`. Despite the popularity of the medical domain for users of search engines, and current interest in this topic within the Information Retrieval research community, development of search and access technologies remains particularly challenging. This workshop brought together researchers working on this challenge, in particular those working at the interface between medical text and image retrieval. The results of the VISCERAL Retrieval Benchmark were also presented at this workshop. The proceedings are available in Springer LNCS [5].

7 Student Support

Due to funding from the ELIAS Network of the European Science Foundation and the British Computer Society Information Retrieval Specialist Group, the ECIR 2015 was able to offer significant student support. Based on a competitive selection from 33 student grant applications, 22 students were awarded a registration fee waiver, and 19 of these students were also awarded a contribution toward their travel and accommodation. These students provided extensive assistance to the organisers by carrying out tasks to support the running of the conference.

8 Conclusions

There were two main innovations at the ECIR 2015. The first was the Reproducible IR Research Track, which encouraged some very interesting papers, presentations and discussion. In particular, the discussion among the members of the Programme Committee for this track was intense, leading to the identification of the following three aspects of reproducibility:

Scientific aspect: an experimental result is independently reproduced under different conditions

Impact aspect: a surprising result is reproduced and shown to hold, or a result obtained on proprietary data is shown to hold on publicly available data

Further reproducibility aspect: All materials (code and data) used to reproduce the result are made available

Only the first aspect was explicitly requested from papers submitted to the Reproducible IR track in 2015, and therefore acceptance decisions were made based only on this aspect. Future tracks could consider expanding the requirements to the second and third aspect.

The second innovation was the new format of the Industry Day, focussing more on start-ups and niche companies. The latter was a great success, especially in terms of the amount of audience participation. A useful suggestion from the attendees for future ECIR conferences is to have a structured way of putting together the companies looking to hire people with those people looking for jobs.

The next ECIR[1] will be held in Padua, Italy, from the 20th to the 23rd of March 2016.

9 Acknowledgements

We thank the organising committee of the ECIR 2015: Industry Day Chairs – Jussi Karlgren (Gavagai and KTH, Sweden) and Paul Ogilvie (LinkedIn, USA); Tutorial Chairs – Birger Larsen (Aalborg University, Denmark) and Adrian Iftene ("Al.I.Cuza" University Iasi, Romania); Workshop Chairs – Guido Zuccon (Queensland University of Technology, Australia), András Benczur (Hungarian Academy of Sciences, Hungary); Demo Chairs – João Magalhães (Universidade Nova de Lisboa, Portugal) and Jan Šnajder (University of Zagreb, Croatia); Student Mentor Chairs – Pia Borlund (Royal School of Library and Inf. Sci., Denmark) and Michal Laclavík (Slovak Academy of Sciences, Slovakia); and Publicity Chair – Ralf Bierig (TU Wien, Austria). We also thank the local organisers: Linda Andersson, Aldo Lipani, Mihai Lupu, João Palotti, Florina Piroi, Navid Rekabsaz, Serwah Sabetghadam and Veronika Stefanov.

Photographs were taken by Aldo Lipani.

The workshop and tutorial descriptions were adapted from the abstracts written by the workshop and tutorial organisers.

Finally, ECIR 2015 would not have been possible without the generous financial support from our sponsors: Google (gold level); Yahoo! Labs and Yandex (silver level); Precognox and max.recall (bronze level). The conference was supported by the Information Retrieval Specialist Group at the British Computer Society (BCS-IRSG), the ELIAS Research Network Program of the European Science Foundation, and the City of Vienna.

[1] `http://ecir2016.dei.unipd.it/`

References

[1] Maria Gäde, Mark Michael Hall, Hugo Huurdeman, Jaap Kamps, Marijn Koolen, Mette Skov, Elaine Toms, and David Walsh, editors. *SCST 2015: Supporting Complex Search Tasks*, volume 1338. CEUR Workshop Proceedings, 2015.

[2] Allan Hanbury, Gabriella Kazai, Andreas Rauber, and Norbert Fuhr, editors. *Advances in Information Retrieval*, volume 9022 of *LNCS*. Springer, 2015.

[3] Frank Hopfgartner, Gabriella Kazai, Udo Kruschwitz, and Michael Meder, editors. *GamifIR 2015: Second International Workshop on Gamification for Information Retrieval*, volume 1345. CEUR Workshop Proceedings, 2015.

[4] Philipp Mayr, Ingo Frommholz, and Peter Mutschke, editors. *BIR 2015: Bibliometric-enhanced Information Retrieval*, volume 1344. CEUR Workshop Proceedings, 2015.

[5] Henning Müller, Oscar Jiménez del Toro, Allan Hanbury, Georg Langs, and Antonio Foncubierta Rodríguez, editors. *Multimodal Retrieval in the Medical Domain*, volume 9059 of *LNCS*. Springer, 2015.