ECIR 2010: 32nd European Conference on Information Retrieval Research

Cathal Gurrin
School of Computing
Dublin City University, Ireland
cgurrin@computing.dcu.ie

Yulan He
Knowledge Media Institute
The Open University, UK
y.he@open.ac.uk

Udo Kruschwitz
School of Computer Science and Electronic Engineering
University of Essex, UK
udo@essex.ac.uk

Suzanne Little and Stefan Rüger
Knowledge Media Institute
The Open University, UK
{*s.little, s.rueger*}*@open.ac.uk*

1 Overview

The 32nd European Conference on Information Retrieval (ECIR 2010)[1] was held at The Open University during March 28-31, 2010 in Milton Keynes, United Kingdom. The conference was organized by the Knowledge Media Institute (KMi), The Open University, in co-operation with Dublin City University and the University of Essex. The core conference attracted well over 160 attendees, and about 210 individuals attended one or more events of the overall programme that included tutorials and workshops before the conference and the Industry Day after the conference.

The ECIR conference series is the annual conference of the Information Retrieval Specialist Group of BCS, The Chartered Institute for IT, (BCS IRSG) and had been running in various forms since 1979. ECIR has long been the foremost Information Retrieval (IR) conference in Europe. It started out as a small-scale colloquium, but increased in size until it became a conference in 2003. Irrespective of this change, the conference has kept the ethos of encouraging the participation of research students or new researchers to the area either by paper or poster presentation or by supporting attendance at the conference. ECIR is a general and multi-disciplinary information retrieval conference that accepts submissions on all aspects of IR including those from computer scientists, information scientists, psychologists, engineers etc. Since 1998 the conference has alternated between the UK and mainland

[1] http://kmi.open.ac.uk/events/ecir2010

Europe: recent instances of the conference include Toulouse (2009), Glasgow (2008), Rome (2007), London (2006), Santiago de Compostela (2005), Sunderland (2004), Pisa (2003) and Glasgow (2002).

202 full-paper submissions were reviewed by at least 3 members of an International Programme committee, and 44 full papers were eventually accepted for presentation at ECIR. The ensuing acceptance rate of 22% (44/202) highlights what a competitive conference ECIR has become. In line with previous ECIR conferences, a large proportion of submissions actually came from outside Europe. Full papers were submitted from Continental Europe (40%), UK (14%), North and South America (15%), Asia and Australia (28%), Middle East and Africa (3%).

ECIR has traditionally been a conference with a strong student focus, which is evidenced by the fact that 61% of the accepted full papers had a student as first author. To allow as much interaction between delegates as possible and to keep in the spirit of the conference, it was decided to run ECIR 2010 as a single-track event. As a result, ECIR 2010 introduced two presentation formats for full papers; some were presented orally, others in small groups as poster format during lunch time. The presentation format did not signify any difference in quality, rather the presentation format was decided after the full papers had been accepted at the Program Committee meeting held at the University of Essex. The views of the reviewers and the nature of the paper was taken into consideration when selecting the most appropriate presentation format for each paper.

The total of 72 accepted papers, posters and demonstrations consisted of 35 from Continental Europe (49%), 14 from UK (19%), 16 from North and South America (22%), 7 from Asia and Australia (10%). The accepted contributions represent the state of the art in information retrieval (IR) and cover a diverse range of topics including NLP and text mining, Web IR, evaluation, multimedia IR, distributed IR and performance issues, IR theory and formal models, personalization and recommendation, domain-specific IR, cross-language IR, and user issues.

In addition to the full research papers, ECIR 2010 also accepted poster and demonstration submissions (demonstration papers were newly introduced this year). In total, 23 posters and 5 demonstrations were accepted for ECIR 2010.

2 Keynote Talks

The conference featured three keynote talks:

- Barry Smyth (University College Dublin) on *Web Search Futures: Personal, Collaborative, Social*
- Hinrich Schütze (University of Stuttgart) on *Natural Language Processing in IR - Shallow or Deep?*
- Mirella Lapata (University of Edinburgh) on *Image and Natural Language Processing for Multimedia Information Retrieval*

Barry Smyth illustrated how Web search is changing and how much social networking has already impacted on our everyday use of the Web. He described the future of Web search as

being much more collaborative in nature. The benefit of this social aspect should be clear in a world where two out of three searches submitted by a user are for something that a friend has found recently. He expects the future of search also to be more contextualised and more adaptive with a move from relevance to reputation.

Hinrich Schütze looked at the future of Web search from a different angle. He started off by outlining how search in the "real world" is not a one-shot query but typically an interactive dialogue, but today's search engines are only good for one-shot queries. Dialogues will require more natural language processing (NLP) and understanding. NLP is already being used in mainstream search applications but a lot of it is shallow, e.g. as used in question-answering, clarification dialogues, dynamic summaries etc. Deep NLP has been used but so far there seems to be no killer application yet for parsing, for example. Unlike relevance feedback (which is based on keywords), interactive IR offers an opportunity for deep linguistic processing. Visualization will be key to progress along those lines as interactive IR can otherwise be too much work for the user. Hinrich sees the main enabler for deep NLP in robust speech processing as speaking a long query is very little effort for a user.

The third keynote talk was allocated for the winner of the 2009 Karen Spärck Jones Award. This award was created by the BCS IRSG in conjunction with the BCS to commemorate the achievements of Karen Spärck Jones (KSJ). 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 2009 KSJ Award was given to Mirella Lapata. She was expected to give birth shortly after the conference and, on the advice of two medical consultants, refrained from travelling to ECIR. Instead, she shared her results of work in image and natural language processing for multimedia information retrieval via a pre-recorded video presentation. In her presentation she argued that the solution to multimedia search lies in combining both image and text processing. Questions from the audience were then noted and passed back to Mirella.

3 Conference Programme

ECIR 2010 initially received a total of 217 full-paper submissions, 15 of which were withdrawn or did not to conform to submission guidelines (for example, were not written in English). Hence, a total 202 papers were dispatched to reviewers drawn from an international Program Committee. Each of these papers was reviewed by at least three experts, and 44 were selected in a *triple-blind process* as full research papers. "Triple blind" means submissions had to be anonymous, reviewer identities were kept anonymous to both the authors *and* the participants of the Programme Committee (PC) meeting in November at the University of Essex. Every reviewer had been invited to attend the PC meeting in person, and arrangements were made for people to phone in if they could not attend in person, but wanted to be "present". At the meeting each submission was discussed as to its merits and weaknesses based on the underlying written justification of the reviews (rather than solely ranked by marks). In rare cases where the justification of the given marks was weak, incomprehensible or contradictive, a further on-the-spot review carried out. The organisers of the conference did not submit papers of their own, and PC participants had to leave the room at the PC meeting when a paper was discussed where one of the authors shared an affiliation with them. This format kept the discussions focussed on the quality of the submission and

made sure that scientific quality and novelty were the overriding criteria for acceptance.

This section briefly introduces full papers arranged according to their corresponding sessions. It needs of course be pointed out that deciding on any such grouping can be an art in its own right.

3.1 NLP and Text Mining

The paper entitled *A Language Modeling Approach for Temporal Information Needs* by *K. Berberich, S. Bedathur, O. Alonso, and G. Weikum* tackles the inherent uncertainty problem in temporal expressions in users' queries by integrating temporal expressions into a language modeling and thereby making them first-class citizens of the retrieval model and considering their inherent uncertainty.

J. Martinez-Romo and L. Araujo in their paper: *Analyzing Information Retrieval Methods to Recover Broken Web Links* compare different techniques to automatically find candidate web pages to substitute broken links. They extract information from the anchor text, the content of the page containing the link, and the cache page in some digital library. Term frequencies and a language modeling approach are used to select terms to construct the queries submitted to a search engine, while co-occurrence measures and a language model approach are applied for ranking the final results.

The paper by *J. Karlgren, G. Eriksson, M. Sahlgren, and O. Täckström* entitled *Between Bags and Trees — Constructional Patterns in Text Used for Attitude Identification* proposes to use non-terminological information, structural features of text, in addition to the presence of content terms, to find attitudinal expressions in written English text. They find that constructional features transfer well across different text collections.

F. Boudin, L. Shi, and J.-Y. Nie in their paper entitled *Improving Medical Information Retrieval with PICO Element Detection* tackle the problem of identifying appropriate resources and searching for the best available evidence for medical treatment in evidence-based medicine (EBM) by exploring the incorporation of the PICO elements, Population/Problem (P), Intervention (I), Comparison (C) and Outcome (O), into the IR process. They found that the I and P elements can be used to enhance the retrieval process and thereby give significantly better retrieval effectiveness than the state-of-the-art methods.

The paper entitled *The Role of Query Sessions in Extracting Instance Attributes from Web Search Queries* by *M. Pasca, E. Alfonseca, E. Robledo-Arnuncio, R. Martin-Brualla, and K. Hall* proposes a weakly-supervised extraction method which exploits anonymized Web-search query sessions, as an alternative to isolated, individual queries, to acquire per-instance attributes (e.g., *top speed* for *chevrolet corvette*, or *population density* for *Brazil*). Inherent challenges associated with using sessions for attribute extraction, such as a large majority of within-session queries not being related to attributes, are overcome by using attributes globally extracted from isolated queries as an unsupervised filtering mechanism.

The paper entitled *Transliteration Equivalence Using Canonical Correlation Analysis* by *R. Udupa and M.M. Khapra* addresses the problem of Transliteration Equivalence, i.e. determining whether a pair of words in two different languages are name transliterations or not, by considering name transliterations in two languages as two views of the same semantic object and compute a low-dimensional common semantic feature space using Canonical Correlation Analysis (CCA). Similarity of the words in the common semantic feature space

forms the basis for classifying a pair of names as transliterations.

3.2 Web IR

R.L.T. Santos, J. Peng, C. Macdonald, and I. Ounis in their paper entitled *Explicit Search Result Diversification through Sub-queries* introduce xQuAD, a novel model for search result diversification that attempts to build such a diversified ranking by explicitly accounting for the relationship between documents retrieved for the original query and the possible aspects underlying this query, in the form of sub-queries.

S. Stamou and E.N. Efthimiadis in their paper: *Interpreting User Inactivity on Search Results* study users' inactivity on search results in relation to their pursued search goals and investigate the impact of displayed results on user clicking decisions. Their study examines two types of post-query user inactivity, pre-determined and post-determined, depending on whether the user started searching with a preset intention to look for answers only within the result snippets and did not intend to click through the results, or the user inactivity was decided after the user had reviewed the list of retrieved documents.

The paper entitled *Learning to Select a Ranking Function* by *J. Peng, C. Macdonald, and I. Ounis* proposes a novel Learning To Select framework that selectively applies an appropriate ranking function on a per-query basis. The use of divergence, which measures the extent that a document ranking function alters the scores of an initial ranking of documents for a given query, is employed as a query feature to identify similar training queries for an unseen query. The ranking function which performs the best on this identified training query set is then chosen for the unseen query.

The paper entitled *Mining Anchor Text Trends for Retrieval* by *N. Dai and B.D. Davison* argues that historical trends of anchor text importance is important to web search. They propose a novel temporal anchor text weighting method to incorporate the trends of anchor text creation over time, which combines historical weights on anchor text by propagating the anchor text weights among snapshots over the time axis.

The paper by *K. Collins-Thompson and P.N. Bennett* entitled *Predicting Query Performance via Classification* investigates using topic prediction data, as a summary of document content, to compute measures of search result quality. Their findings suggest that class-based statistics can be computed efficiently online and using class predictions can offer comparable performance to full language models while reducing computation overhead.

3.3 Distributed IR and Performance Issues

I. Cox, J. Zhu, R. Fu, and L.K. Hansen in their paper entitled *Approximately Correct (PAC) Search* propose a modification to a non-deterministic architecture for IR, the PAC architecture, by introducing a centralized query coordination node. To respond to a query, random sampling of computers is replaced with pseudo-random sampling using the query as a seed. Then, for queries that occur frequently, this pseudo-random sample is iteratively refined so that performance improves with each iteration. Experiments on the TREC-8 dataset demonstrate that for queries that occur 10 or more times, the performance of a non-deterministic PAC architecture can closely match that of a deterministic system.

The paper entitled *Learning to Distribute Queries into Web Search Nodes* by *M. Mendoza, M. Marín, F. Ferrarotti, and B. Poblete* investigates machine learning algorithms to distribute queries onto web search nodes and propose a logistic regression model to quickly predict the most pertinent search nodes for a given query.

The paper by *O. Papapetrou, W. Siberski, and N. Fuhr* entitled *Text Clustering for Peer-to-Peer Networks with Probabilistic Guarantees* presents a text clustering algorithm for peer-to-peer networks with high scalability by using a probabilistic approach for assigning documents to clusters. It enables a peer to compare each of its documents only with very few selected clusters, without significant loss of clustering quality.

The paper entitled *XML Retrieval Using Pruned Element-Index Files* by *I.S. Altingovde, D. Atilgan, and Ö. Ulusoy* proposes using static index pruning techniques for obtaining more compact index files that can still result in comparable retrieval performance to that of a full index in XML retrieval.

3.4 Multimedia IR

K. Athanasakos, V. Stathopoulos, and J.M. Jose introduce a framework for the evaluation of image annotation models in their paper entitled *A Framework for Evaluating Automatic Image Annotation Algorithms*, which they use to evaluate two state-of-the-art Automatic Image Annotation (AIA) algorithms. They reveal that a simple SVM approach using Global MPEG-7 Features outperforms state-of-the-art AIA models across several collection settings.

H.-S. Kim, H.-W. Chang, J. Lee, and D. Lee in their paper entitled *BASIL: Effective Near-Duplicate Image Detection Using Gene Sequence Alignment* propose a new algorithm, termed as "BlASted Image Linkage" (BASIL), which uses the popular gene sequence alignment algorithm BLAST in Biology in detecting near-duplicate images. They study how various image features and gene sequence generation methods (using gene alphabets such as A, C, G, and T in DNA sequences) affect the accuracy and performance of detecting near-duplicate images.

R. Aly, A. Doherty, D. Hiemstra, and A. Smeaton propose a retrieval framework for news story items consisting of multiple shots in their paper entitled *Beyond Shot Retrieval: Searching for Broadcast News Items Using Language Models of Concepts*. The framework consists of a concept based language model which ranks news items with known occurrences of semantic concepts by the probability that an important concept is produced from the concept distribution of the news item and a probabilistic model of the uncertain presence, or risk, of these concepts.

The paper by *S. Peña Saldarriaga, E. Morin, and C. Viard-Gaudin* entitled *Ranking Fusion Methods Applied to On-Line Handwriting Information Retrieval* presents an empirical study on the application of ranking fusion methods in the context of handwriting information retrieval. Retrieval approaches on texts obtained through handwriting recognition and recognition-free methods using word-spotting algorithms are combined to give a significant effectiveness improvement.

3.5 IR Theory and Formal Models

K. Balog, M. Bron, and M. de Rijke present a general probabilistic framework for category-based query modeling for entity search in their paper entitled *Category-Based Query Modeling for Entity Search*. They focus on the use of category information and show the advantage of a category-based representation over a term-based representation, and also demonstrate the effectiveness of category-based expansion using example entities.

The paper by *S. Agarwal and M. Collins* entitled *Maximum Margin Ranking Algorithms for Information Retrieval* proposes a family of ranking algorithms by optimizing variations of the hinge loss used in support vector machines. The algorithms preserve the simplicity of standard pair-wise ranking methods in machine learning, yet show performance comparable to state-of-the-art IR ranking algorithms.

W. Zheng and H. Fang in their paper: *Query Aspect Based Term Weighting Regularization in Information Retrieval* study the incorporation of query term relations into existing retrieval models by first developing a general strategy that can systematically integrate a term weighting regularization function into existing retrieval functions, and then proposing two specific regularization functions based on the guidance provided by constraint analysis.

The paper entitled *Using the Quantum Probability Ranking Principle to Rank Interdependent Documents* by *G. Zuccon and L. Azzopardi* explores whether the new Quantum Probability Ranking Principle (QPRP), which implicitly captures dependencies between documents through "quantum interference", leads to improved performance for subtopic retrieval, where novelty and diversity is required. They show that QPRP consistently outperforms the previous ranking strategies such as Probability Ranking Principle, Maximal Marginal Relevance and Portfolio Theory.

X. Tu, T. He, L. Chen, J. Luo, and M. Zhang in their paper entitled *Wikipedia-Based Semantic Smoothing for the Language Modeling Approach to Information Retrieval* present a novel Wikipedia-based semantic smoothing method that decomposes a document into a set of weighted Wikipedia concepts and then maps those unambiguous Wikipedia concepts into query terms. The mapping probabilities from each Wikipedia concept to individual terms are estimated through the EM algorithm. Document models based on Wikipedia concept mapping are then derived.

3.6 User Issues

L. Kelly and G.J.F. Jones presents a novel query independent static biometric scoring approach for re-ranking result lists retrieved from a lifelog in their paper entitled *Biometric Response as a Source of Query Independent Scoring in Lifelog Retrieval*, which contains digital records captured from an individual's daily life, e.g. emails, web pages downloaded and SMSs sent or received, using a BM25 model for content and content + context data. They explored the utility of galvanic skin response (GSR) and skin temperature (ST) associated with past access to items as a measure of potential future significance of items.

N. Gooda Sahib, A. Tombros, and I. Ruthven in their paper entitled *Enabling Interactive Query Expansion through Eliciting the Potential Effect of Expansion Terms* investigate a method of increasing the uptake of Interactive Query Expansion by displaying summary previews that allow searchers to view the impact of their expansion decisions in real time, engage more with suggested terms, and support them in making good expansion decision.

The paper by *S. Kriewel and N. Fuhr* entitled *Evaluation of an Adaptive Search Suggestion System* describes an adaptive search suggestion system based on case-based reasoning techniques, and details an evaluation of its usefulness in helping users employ better search strategies.

R. Kaptein, D. Hiemstra, and J. Kamps in their paper entitled *How Different Are Language Models and Word Clouds?* investigate what tag clouds and language modeling approaches can learn from each other, and specifically whether language modeling techniques can be used to generate tag or word clouds automatically from documents.

3.7 Personalization and Recommendation

The paper entitled *A Performance Prediction Approach to Enhance Collaborative Filtering Performance* by *A. Bellogín and P. Castells* investigates the adaptation of clarity-based query performance predictors to define predictors of neighbor performance in Collaborative Filtering (CF). A predictor is proposed and introduced in a memory-based CF algorithm to produce a dynamic variant where neighbor ratings are weighted based on their predicted performance.

J. Redpath, D.H. Glass, S. McClean, and L. Chen in their paper: *Collaborative Filtering: The Aim of Recommender Systems and the Significance of User Ratings* investigate the significance of user ratings in recommender systems by considering their inclusion/exclusion in both the generation and evaluation of recommendations. They argue that it is important to identify the aim of the system before evaluating the accuracy of a recommender algorithm since the use of ratings would generate different results when different evaluation metrics are used.

The paper by *T. Jambor and J. Wang* entitled *Goal-Driven Collaborative Filtering — A Directional Error Based Approach* argues that defining an error function that is uniform across rating scales in collaborative filtering is limited since different applications may have different recommendation goals and thus error functions. They propose a flexible optimization framework that can adapt to individual recommendation goals by introducing a weight function to capture the cost (risk) of each individual predictions which can be learned from the specified performance measures.

The paper entitled *Personalizing Web Search with Folksonomy-Based User and Document Profiles* by *D. Vallet, I. Cantador, and J.M. Jose* proposes to represent a user profile in terms of social tags, manually provided by users in folksonomy systems to describe, categorize and organize items of interest, and investigate a number of novel techniques that exploit the user's social tags to re-rank results obtained with a Web search engine.

M. Harvey, M. Baillie, I. Ruthven, and M.J. Carman in their paper entitled *Tripartite Hidden Topic Models for Personalised Tag Suggestion* extend the latent Dirichlet allocation topic model to include user data and use the estimated probability distributions to provide personalized tag suggestions to users.

3.8 Domain-Specific IR and CLIR

The paper entitled *Extracting Multilingual Topics from Unaligned Comparable Corpora* by *J. Jagarlamudi and H. Daumé III* presents a generative model called JointLDA which uses

a bilingual dictionary to mine multilingual topics from an unaligned corpus. They speculate that the JointLDA model has better predictive power compared to the bag-of-word based translation model leaving the possibility for JointLDA to be preferred over bag-of-word model for cross-lingual IR applications.

S. Bashir and A. Rauber in their paper entitled *Improving Retrievability of Patents in Prior-Art Search* expand prior-art queries generated from query patents using query expansion with pseudo relevance feedback in patent retrieval. They propose a novel approach to automatically select better terms from query patents based on their proximity distribution with prior-art queries that are used as features for computing similarity measures.

L. Shi in his paper: *Mining OOV Translations from Mixed-Language Web Pages for Cross Language Information Retrieval* presents a method that automatically acquires a large quantity of OOV translations from the web by adaptively learning translation extraction patterns based on the observation that translation pairs on the same page tend to appear following similar layout patterns.

The paper by *J. Soo and O. Frieder* entitled *On Foreign Name Search* proposes a language-independent foreign names search approach, called Segments, and compares it against traditional n-gram and Soundex based solutions.

X. Yin, X. Huang, and Z. Li in their paper entitled *Promoting Ranking Diversity for Biomedical Information Retrieval Using Wikipedia* propose a cost-based re-ranking method to promote ranking diversity for biomedical information retrieval which aims at finding passages that cover many different aspects of a query topic.

R. Schenkel in his paper: *Temporal Shingling for Version Identification in Web Archives* presents temporal shingling, an extension of the well-established shingling technique for measuring how similar two snapshots of a page are, for version identification in web archives. The method considers the lifespan of shingles to differentiate between important updates that should be archived and transient changes that may be ignored.

3.9 Evaluation

C. Hauff, D. Hiemstra, L. Azzopardi, and F. de Jong in their paper entitled *A Case for Automatic System Evaluation* perform a wider analysis of system ranking estimation methods on 16 TREC data sets which cover more tasks and corpora than previously tested. Their analysis reveals that the performance of system ranking estimation approaches varies across topics and can be improved by selecting the "right" subset of topics from a topic set. They also observe that the commonly experienced problem of underestimating the performance of the best systems is data set dependent and not inherent to system ranking estimation methods.

The paper by *H.D. Kim, C. Zhai, and J. Han* entitled *Aggregation of Multiple Judgments for Evaluating Ordered Lists* proposes three new methods for aggregating multiple order judgments to evaluate ordered lists; weighted aggregation, combined ranking, frequent sequential pattern based aggregation. The first method assigns different weights on experts judgments based on the consensus and compute evaluation score based on weights. The second method finds combined ranking from multiple judgments and uses it for evaluation. The third method collects frequent sequential patterns and their frequencies from all the human judgments of ordering and scores an ordered list based on how well the list matches

these frequent sequential patterns.

J. Tang and M. Sanderson conduct evaluation and user preference study on spatial diversity in the context of spatial information retrieval in their paper entitled *Evaluation and User Preference Study on Spatial Diversity*. They show the potentials of spatial diversity by not only the traditional evaluation metrics (precision and cluster recall), but also through a user preference study using Amazon Mechanical Turk.

M. Tsagkias, W. Weerkamp, and M. de Rijke in their paper entitled *News Comments: Exploring, Modeling, and Online Prediction* explore the news comments space, and compare the log-normal and the negative binomial distributions for modeling comments from various news agents. They also examine the feasibility of predicting the number of comments, based on the number of comments observed shortly after publication.

C. Hauff, L. Azzopardi, D. Hiemstra, and F. de Jong in their paper entitled *Query Performance Prediction: Evaluation Contrasted with Effectiveness* control the quality of a query performance predictor in order to examine how the strength of the correlation achieved affects the effectiveness of an adaptive retrieval system. They show that many existing predictors fail to achieve a correlation strong enough to reliably improve the retrieval effectiveness in the Selective Query Expansion as well as the Meta-Search setting.

3.10 Posters & Demonstrations

ECIR 2010 received a total of 73 poster/demonstration submissions, each of which were reviewed by at least three members of the Program Committee. Out of these submissions, 23 posters and 5 demonstrations were accepted. The poster and demonstration session took place on the first evening of ECIR 2010 during the conference reception. The posters presented at ECIR 2010 spanned a full range of topics of interest to the IR community, with a notable interest in social and Web IR, cross-language IR and ranking models. The poster session stimulated much discussion at the conference and really acted as a good interaction and focal point. The demonstrations also took place during the same session. Even though this was the first time that ECIR hosted demonstration papers, the five excellent demonstrations attracted keen interest from the gathered participants and the demostration papers and associated session could be considered a success at ECIR 2010.

4 Awards

In the same reception that hosted the posters and demonstrations, the award announcements were made. Three committees comprising three ECIR 2010 PC members each were set up to select the best paper, best student paper and the best poster. No best student paper award was given out as the two joint best papers had both a student as its prime author:

- *Evaluation of an adaptive search suggestion system* by Sascha Kriewel and Norbert Fuhr (University of Duisburg-Essen)
- *Promoting Ranking Diversity for Biomedical Information Retrieval Using Wikipedia* by Xiaoshi Yin, Jimmy Huang and Zhoujun Li (Beihang University Beijing and York University Toronto).

The best poster award was given to *Filtering document with subspaces* by Benjamin Piwowarski, Ingo Frommholz, Yashar Moshfeghi, Mounia Lalmas, Keith van Rijsbergen (University of Glasgow).

5 Workshops & Tutorials

As in the previous two years, ECIR was preceded by a day of workshops and tutorials. This format has proven to be very successful; at ECIR 2010 there were four tutorials and four workshops. To encourage participation in large numbers the registration fee for workshops and tutorials was waived for those attendees who registered for the main conference. This has paid off as the events were very well attended. The accepted tutorials and workshops offered a diverse range of topics and the tutorials in particular involved a significant element of practical work. For example, the user interface tutorial — delivered by Tony Russell-Rose (Endeca) — required participants to work in groups to analyze a highly popular but poorly designed Web site and go beyond identifying the problems and outline solutions. The crowdsourcing tutorial by Omar Alonso (Microsoft Bing) illustrated the power of the wisdom of the crowd through live experiments run on Amazon's Mechanical Turk platform.

6 Industry Day

ECIR 2010 was followed by an Industry day on Thursday 1st April 2010 devoted to the interests and needs of Information Retrieval practitioners – ECIR Industry Day[2]. In common with previous BCS IRSG Industry Days, the day represented a unique opportunity to meet and share ideas with a broad mix of experts from both industry and academia. The speakers presented the latest challenges faced by developers and practitioners in both enterprise and web search. This was the third ECIR Industry Day following successful events at ECIR 2006 in London and ECIR 2008 in Glasgow. This year, Industry Day attracted a record number of more than 70 participants (about a third of which also attended the main conference).

A specific feature of ECIR Industry Day is that the talks are given by experts in the area *and not* sales representatives. This ensures that participants get a feel for what the underlying technical issues are in state-of-the-art search solutions and applications. This year, one of the common themes was applications of semantic analysis, a rapidly maturing academic field, now ready for wider adoption in the marketplace.

Talks were grouped into four different sessions. The day started with *Advancing Web Search*, a session that brought together speakers from Yahoo (Ricardo Baeza-Yates), Google (Dan Crow) and Microsoft Bing (Milad Shokouhi), all of them heavy-weights in the IR community. All three talks nicely illustrated how far the field have moved away from simple bag-of-words approaches towards "semantic search".

The second session featured three very different views on *Enterprise Search*, namely that of a user interface designer (Vegard Sandvold of Comperio), the provider of an enterprise search product with some very practical suggestions (David Hawking of Funnelback), and that of an analyst (Nick Patience of The 451 Group).

[2] http://csee.essex.ac.uk/staff/udo/ecir2010

The first session after lunch (*Domain-Specific Search*) started with a talk by Rob Blackwell (Active Web Solutions) demonstrating that search is more than just Web search, in particular if you are working in a heavily constrained environment where facts (rather than documents) are required and high recall might be important. Sally Chambers (The European Library) voiced the view of digital library providers who require search products that can deal with heavily structured but also unstructured data. She pointed out that Web search progress has made the life of librarians more difficult as users want a search engine just like that when they search library catalogues. The final talk in this session was given by John Tait (Information Retrieval Facility) who promoted collaborative research, technology transfer and networking and illustrated these points using the IRF Innovation Cycle.

In the final session, *Beyond Search: Question-Answering*, the speakers gave us a glimpse into the future of open-domain question-answering. Jon McLoone (Wolfram Alpha) demonstrated the types of questions Wolfram Alpha is capable of answering by accessing a structured knowledge base of quality-checked facts. Simon Overell (True Knowledge) started with a similar motivation but the methods are very different: instead of filling the database with facts provided by in-house experts, True Knowledge uses a core ontology but then relies very much on input provided by the Web and its users. Interestingly, both Jon and Simon agreed that if their QA systems do not understand the user's question, it will rather keep quiet than guessing the right interpretation.

To round up the day, the talks were followed by a panel session. A panel of experts, Dan Crow (Google), Gjergji Kasneci (Microsoft Research), Jon McLoone (Wolfram Alpha), Simon Overell (True Knowledge), discussed their views on *Leveraging semantics to enable better search experiences*. The panel was moderated by Mike Taylor (Microsoft Research). The main discussion revolved around the idea of either using the Web to extract knowledge or to have curated data, two approaches at different ends of the spectrum. The Google approach appears to be very much looking at whatever can be found on the Web and then making sense of that (without trying to be an authoritative resource itself). Quite the opposite approach is taken by Wolfram Alpha and True Knowledge in that they rely on vetted facts and knowledge.

The panel could have continued for much longer, but the white wine was getting warm and so the official programme was closed and attendees kept discussing a variety of issues over drinks and nibbles. Despite best efforts we were unable to fully finish the wine supply. Nevertheless, the format of Industry Day has been shown to be highly successful as part of ECIR.

7 Social Programme

Conferences are a great way to meet colleagues and discuss emerging trends. ECIR provided more than 8.5 hours of coffee and lunch breaks over the three days as well as a reception on the first day and a banquet on the second day of the conference. None of these were meant for idle conversation: The reception, which carried on well beyond the advertised finish of 8pm, hosted the poster and demonstration session, while the lunch breaks provided the backdrop for the 18 full papers that were presented in poster format.

Prior to each session, participants could play the highly interactive 3D BallBouncer game: A live video of the audience is shown as a mirror image on screen into which virtual beach

Figure 1: Say "Information Retrieval": participants at the banquet venue Bletchley Park, where Alan Turing guided naval cryptanalysis during the second World War. Photograph © by Chris Valentine, used with kind permission. More Pictures on flickr: http://www.flickr.com/search/?q=ecir2010.

balls are thrown; the real audience could then very realistically hit the virtual balls, which reacted in real-time according to the movements of the audience (even more fun when a line on the screen split the audience into two competing teams while the beach balls turned into ticking time bombs that each team had to make sure to detonate in the other half).

The ECIR 2010 Conference banquet was held at Bletchley Park which has a unique significance for the history of British Computing, being Britain's main decryption site during the Second World War. It was at Bletchley Park that Alan Turing, who later published groundbreaking papers on computability, devised techniques for breaking German ciphers. Turing's wartime work included the design of an electro-mechanical machine nicknamed *bombe* that could find settings for the German's main encryption mechanism, the enigma. The formal sit-down banquet in the Bletchley Park's Mansion house was preceded by a a tour of the grounds, Bletchley Park museum and to the rebuild of the world's first programmable, digital, electronic, computing device in 1944, Colossus, which used vacuum valve tubes to perform the calculations.

Many ECIR participants deployed their skills to try to retrieve messages from a difficult code-breaking puzzle that was hidden inside the banquet menus. While most arrived at plausible conclusion that information retrieval is just so much harder when the sender of a message doesn't want it to be understood by everyone, some managed to crack the messages in time for a prize to be awarded to them at the closing of the conference.

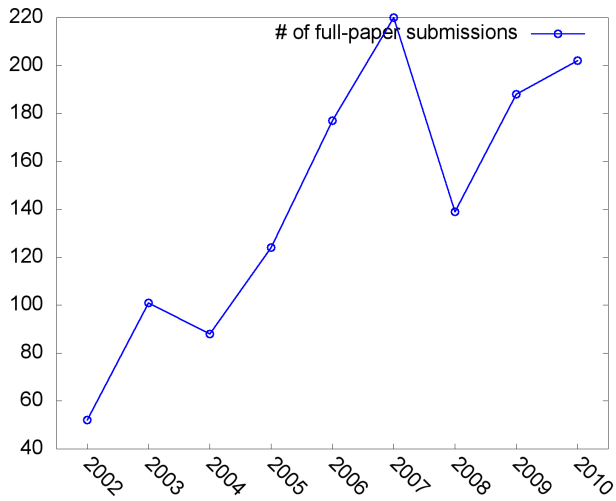

Figure 2: Number of full-paper submissions to ECIR over the years. These numbers were sourced from the prefaces of the corresponding Springer proceedings.

8 ECIR in historic context

ECIR started out as a colloquium in 1979 in Leeds and was held annually at one place or another in the UK until 1998, when the colloquium took place in mainland Europe for the first time. As the number of submissions rose (see Figure 2) over time, the event became more competitive in terms of acceptance rates (see Figure 3). From 2002, the proceedings were published by Springer Verlag in its Lecture Notes in Computer Science Series, and, finally, in 2003, ECIR was renamed from "European Colloquium on Information Retrieval Research" to "European Conference on Information Retrieval" reflecting its nature. Figure 2, which shows the number of full-paper submissions to ECIR from 2002, justifies this decision in hindsight: There has been a overall upward trend for the number of submissions over the last decade.

This upward trend for full-paper submissions is particularly notable as, beginning with 2004, ECIR also started to accept dedicated poster submissions of late breaking results and work in progress (now typically at a level of 70 submissions per year), while some years later workshops were introduced. Both instruments provide additional capacity for presenting research results, as do a number of conferences that have established themselves in the past decade: ACM CIVR, the international Conference on Image and Video Retrieval; ACM MIR, the international conference on Multimedia Information Retrieval (unfortunately sharing the same dates with ECIR in 2010)[3]; BCS ICTIR, the International Conference on the Theory of Information Retrieval research; IRFC, the Information Retrieval Facility Conference; AIRS, the Asia Information Retrieval Symposium and many more.

Although the capacity in terms of the number of accepted full-papers in the core conference has risen over time, it has not risen in line with the submission numbers. Hence, as Figure 3 shows, the acceptance rate of full papers at ECIR has fallen to around 20–25%.

One distinguishing factor of ECIR from the beginning was its focus on, and support

[3]In 2010 MIR left its traditional autumn spot in the calendar only to be consolidated with CIVR into ACM ICMR from 2011

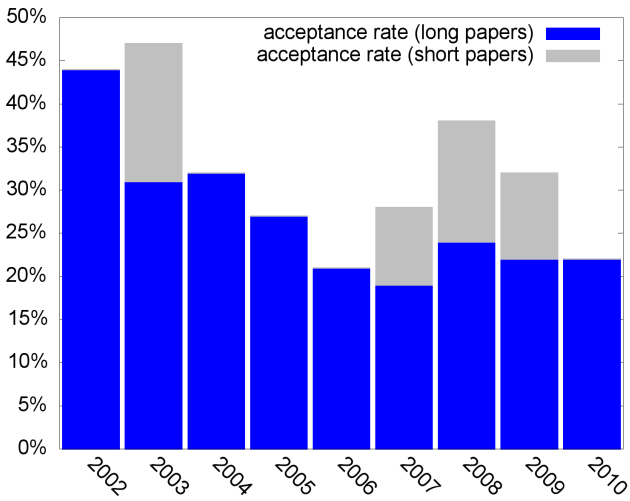

Figure 3: Acceptance rate for (full and short) paper submissions. In some years the programme committee distinguished between long papers and short papers, which then were given fewer pages in the proceedings and a shorter talk time.

for, young researchers. Traditionally ECIR has been a "friendly" environment frequented by senior researchers and new research talent alike. ECIR normally supports student attendance with bursaries (eg, more than 40 in 2006 and 23 in 2010). It turns out that typically the proportion of papers with students as first authors is more than 50% as demonstrated by Figure 4.

9 Feedback and Conclusions

What was new at ECIR 2010? *Tutorials and workshops were free* if one was registered with the main conference. Around 135 participants made use of this, and there was very positive feedback on these satellite activities. *Demonstrations* were introduced this year as well, and they certainly have contributed to the buzz of the conference. ECIR 2010 provided some of the capacity for *full papers through poster delivery format*. Although this is a novel concept for ECIR, it has long been used in other information retrieval conferences such as CIVR and has been deployed in other fields. One difficulty is communicating that the presentation format of a full paper as poster does not signify a "lesser" paper. This is confounded by the fact that ECIR also has established a genuine and different poster stream, which come about by dedicated poster submissions with fewer pages in the proceedings for late-breaking results, significant work in progress or research that is best communicated in an interactive or graphical format. Has the *full paper presented as a poster* in small groups worked out? The feedback that we got was mixed: some liked this format more than parallel sessions, as it gives everyone the opportunity to look at every paper. Others would have preferred parallel sessions, or a longer poster session than the 90 minutes that was allocated to the six posters each day, especially, considering that the session ran during lunchtime.

ECIR 2010 experienced three no-shows from authors and as mentioned, the Karen Spärck Jones award winner presented her keynote through a video with a full-screen slide-show. The no-show authors missed the conference for valid reasons beyond their control. In one case the

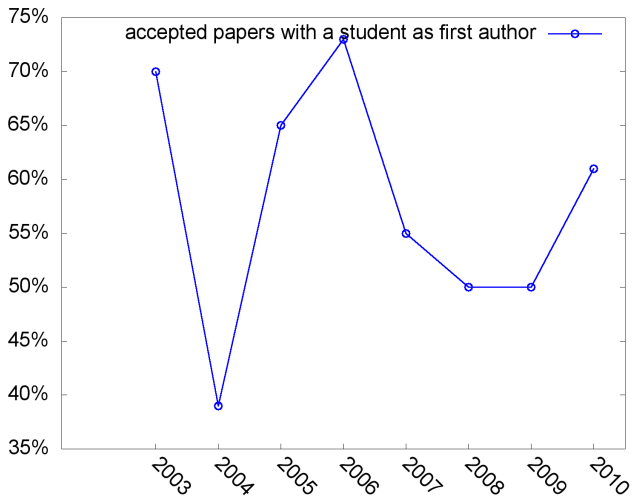

Figure 4: Proportion of accepted full-paper submissions with a student as first author (except for 2007, where this proportion was only available for accepted long papers).

author provided a voice-over slide show to the organisers ahead of time, while in the other cases the authors arranged for a stand-in presenter to present their talk. Looking at the feedback, it was deemed a better solution for authors to provide an alternative presentation, either video or voice-over, rather than having someone else stand in at short notice. Could the organisers not have shifted a full-paper that was scheduled for a poster presentation to give a talk instead of the no-show slot? Not at the short notice available! As organisers we ensured that every paper at the conference had at least someone from the authors registered, but when it only becomes apparent on the presentation day that a particular paper might not have a presenter, there is little that the organisers can do.

Finally, it appears that there are a few buzzwords that seem to arise at any conference on any particular year. In 2010 at ECIR, the notable buzzwords were *diversity* (in the sense of ranking diversity) and (from Industry Day) *faceted search*. In all relevant aspects, however, all feedback agreed: the ECIR 2010 submissions, the reviewing and ultimately the programme were of undoubtedly high quality, the banquet venue was unique and the conference was well worth going to. ECIR 2011[4] will be held in Dublin in April 2011, during the traditional week before Easter. See you there!

Acknowledgements

Our gratitude goes to the ECIR 2010 Program Committee members and other invited reviewers for their objective and thorough reviews of the submitted papers.

We are very grateful to the sponsors for generously providing financial support for the conference including Apriorie Ltd, Active Web Solutions, Exalead, Flax, Funnelback, GCHQ, Google, Information Retrieval Facility, Microsoft Research, and Yahoo Research. In particular, Yahoo Research sponsored the best-paper award, and GCHQ sponsored the banquet in association with Bletchley Park.

[4] http://www.ecir2011.dcu.ie

We would also like to thank the Knowledge Media Institute at The Open University for hosting the event, dealing with all local arrangements, and providing registration and financial management. Our special thanks go to Damian Dadswell (Web), Jane Wilde, Rachel Barnett, Aneta Tumilowicz and The Open University's Finance Division.