



Introduction: Special Issue on the 25th European Conference on Information Retrieval Research

The European Conference on Information Retrieval Research (ECIR) was initially established by the Information Retrieval Specialist Group of the British Computer Society (BCS-IRSG) under the name “Annual Colloquium on Information Retrieval Research”, and always held in the United Kingdom until 1997. Since 1998 the location of the Colloquium has alternated between the United Kingdom and the rest of Europe, in order to reflect the growing European orientation of the event. For the same reason, in 2001 the event was renamed “European Annual Colloquium on Information Retrieval Research”. Since 2002, the proceedings of the Colloquium are being published by Springer Verlag in their Lecture Notes in Computer Science series. As of 2003 BCS-IRSG has decided to rename the event “European Conference on Information Retrieval Research”, in order to reflect what the event had slowly turned into, i.e. a full-blown conference with a European program committee, strong peer reviewing, and a (mostly) European audience. However, ECIR still retains the strong student focus that has characterized the Colloquia from their early inception.

ECIR-03, the 25th “Silver Jubilee” edition, took place in Pisa, Italy, and represented a record-breaking success, both in the number of submitted papers and in the number of participants; a consequence of the former fact was the record-low acceptance rate (31% for full papers), which ensured a high qualitative standard.

This Special Issue brings together the extended versions of seven among the best papers of ECIR-03. The authors of the 10 papers which had been considered the best by the ECIR-03 Program Committee were invited to submit extended versions of their articles. Each of these was reviewed by three to four reviewers, which assessed overall quality of the submitted work and the degree to which this was indeed an extended version of the original ECIR-03 paper. As a result, eight articles were selected for inclusion in this Special Issue. For editorial reasons, one of the selected works (“From Retrieval Status Values to Probabilities of Relevance for Advanced IR Applications”, by Henrik Nottelmann and Norbert Fuhr (University of Duisburg)) had to be included in a previous issue of the journal; however, we may say that it “virtually” belongs to this Special Issue.

The seven papers that appear in this Special Issue range on a variety of topics, from stemming and decomposing, to within-document retrieval, linkage analysis for Web search, recommender systems based on collaborative filtering, language modelling techniques for text classification, topic detection and tracking, and the use of IR technology for e-commerce applications.

The paper “Replicating Web Structure in Small Scale Test Collections”, by Cathal Gurrin and Alan F. Smeaton (Dublin City University) analyzes the reasons why linkage analysis, which is known to be used within several successful Web search engines, has not been shown to bring about systematic benefits in TREC experiments prior to 2002. The

authors track this apparent contradiction to the different linkage patterns which are present in the crawlable portion of the Web, and in current Web test collections, respectively. As a consequence, they identify the link density requirements that test collections for Web search should satisfy in order to better reflect the true linkage patterns of the Web, and describe techniques for building such a collection.

The paper “Within-Document Retrieval: A User-Centred Evaluation of Relevance Profiling”, by David J. Harper, Ivan Koychev, Yixing Sun, and Iain Pirie (The Robert Gordon University) presents both a quantitative and a qualitative analysis of ProfileSkim, a tool that enables users to identify highly relevant passages within long documents already deemed relevant, through the integration of passage retrieval and content-based browsing. The tool is based on the notion of “relevance profiling”, i.e. the computation of how relevance is distributed through the document (which is actually displayed to the user through a tilebar-like interface), and this allows users to effectively locate the passages most relevant to their information needs.

The paper “How Effective is Stemming and Decompounding for German Text Retrieval?”, by Martin Braschler (Eurospider Information Technology AG and Université de Neuchâtel) and Bärbel Ripplinger (Eurospider Information Technology AG), tests a wide range of stemming methods, from empirical to linguistically well-founded ones, on a large collection of long German documents, and using either short or long queries. This experimental setting allows the authors not only to compare different stemming methods with each other (as opposed to comparing the presence with the absence of stemming, which had been the main focus of several previous such studies), but also allows them to obtain results on the effectiveness of these methods when used on collections of long documents and queries, a context in which the improvements due to stemming are less obvious. An analogous study of different compound splitting methods is also presented.

The paper “Augmenting Naive Bayes Classifiers with Statistical Language Models”, by Fuchun Peng (University of Massachusetts at Amherst), Dale Schuurmans and Shaojun Wang (University of Alberta), presents a probabilistic model for text classification in which the standard independence assumption of “naive Bayes” models is relaxed by allowing dependencies among variables to form a Markov chain; this brings about a model intermediate in power between standard naive Bayes models and their “tree augmented” versions. Techniques from statistical n -gram language modelling are used in order to learn these models, and this allows the use of sophisticated smoothing techniques from statistical language modelling for estimating the parameters. The authors experiment this approach on a variety of classification tasks (including genre detection, authorship attribution, and categorization by topical classes) and languages (including ones as diverse as Chinese, English, Greek, and Japanese).

The paper “Simple Semantics in Topic Detection and Tracking”, by Juha Makkonen, Helena Ahonen-Myka and Marko Salmenkivi (University of Helsinki), presents an approach to topic detection and tracking (TDT) in which the authors use an expanded document representation consisting of four different vectors, addressing proper names, locative expressions, temporal expressions, and generic terms, respectively. The first three classes of terms are more semantically determined, and can thus be given a finer-grained interpretation by plugging in class-specific ontologies. This allows the similarity between pieces of texts

representing events to be computed as a combination of class-specific similarity measures, which is helpful since these three classes of terms play a very specific role in TDT.

The paper “OSGS—A Personalized Online Store for E-Commerce Environments”, by Raz Lin, Sarit Kraus (Bar-Ilan University), and Jeffrey Tew (General Motors) applies information retrieval and collaborative filtering technologies to the design of an online store. The authors describe a personalized system for supporting customers who navigate within an online store, using the past history of the user and similar users, plus information on user preference and customer demographics. Their approach provides individual personalized guidance for both browsing and keyword-based searching in an attempt to optimize the tradeoff between finding what the customer needs and finding it quickly.

The paper “Learning User Similarity and Rating Style for Collaborative Recommendation”, by Kwok-Wai Cheung and Lily F. Tian (Hong Kong Baptist University) looks at recommender systems based on collaborative filtering. Within this context, it proposes the use of machine learning techniques for learning the optimal function that measures similarity between users. Furthermore, it proposes a method to reconcile different rating styles adopted by different users, so as to make the ratings obtained from different users comparable and thus correctly interpretable by collaborative filtering algorithms.

Many people have contributed to bringing this Special Issue to life. A special word of thank goes to the external referees, whose job was instrumental in providing timely and high quality feedback to the authors. It is thus a great pleasure to acknowledge the help of Ayse Göker, Barry Smyth, ChengXiang Zhai, David Carmel, Djoerd Hiemstra, Dunja Mladenić, Eero Sormunen, Gabriella Pasi, Gareth Jones, Gianni Amati, Giuseppe Amato, Ian Ruthven, Josiane Mothe, Jussi Karlgren, Keith van Rijsbergen, Lorraine McGinty, Marcello Federico, Martin Braschler, Massimo Melucci, SK Michael Wong, Nicholas Kushmerick, Patrick Gallinari, Peter Ingwersen, Pia Borlund, Renée Pohlmann, Stavros Christodoulakis, Victor Lavrenko, and Wessel Kraaij. I am also grateful to Paul Kantor and Steve Robertson for encouraging me to produce this Special Issue, and to Robbert van Berckelaer and Michelle Misner for their editorial help.

Fabrizio Sebastiani

Istituto di Scienza e Tecnologia dell'Informazione
Consiglio Nazionale delle Ricerche
Pisa, IT