

Preface

The CLEF 2014 conference is the fifteenth edition of the popular CLEF campaign and workshop series which has run since 2000 contributing to the systematic evaluation of multilingual and multimodal information access systems, primarily through experimentation on shared tasks. In 2010 CLEF was launched in a new format, as a conference with research presentations, panels, poster and demo sessions and laboratory evaluation workshops. These are proposed and operated by groups of organizers volunteering their time and effort to define, promote, administrate and run an evaluation activity.

Out of 10 received lab proposals, 8 laboratories have been selected and run during CLEF 2014. To identify the best proposals, besides well-established criteria from previous years' editions of CLEF such as topical relevance, novelty, potential impact on future world affairs, likely number of participants, and the quality of the organizing consortium. This year we further stressed the connection to real-life usage scenarios and we tried to avoid as much as possible overlaps among labs in order to promote synergies and integration. This was possible also thanks to a new activity introduced in 2013, which is a lab organizers and proposers meeting co-located with the European Conference on Information Retrieval (ECIR), held in Moscow on March 26th 2013 and supported by the ELIAS network¹.

CLEF has been always backed by European projects which complemented the incredible amount of volunteering work performed by Lab Organizers and the CLEF community with the resources needed for its necessary central coordination, as it happens for the other major evaluation initiatives as TREC, NTCIR and FIRE. Since 2014, CLEF no longer has support from European projects and it has made an effort to turn itself into a self-sustainable activity. This has been possible thanks to the establishment in late 2013 of the CLEF Association, a no profit legal entity, which, through the support of its members, ensures the resources needed to smoothly run and coordinate CLEF.

Each Lab, building on previous experience, demonstrated maturity coming with new tasks, new and larger data sets, new ways of evaluation or more languages. They are described by the Lab organizers in details, here we just brief on them.

¹ <http://www.elias-network.eu/>

CLEF eHealth - ShARe/CLEF eHealth Evaluation Lab: The usage scenario of the CLEF eHealth lab is to ease patients and next-of-kins' ease in understanding eHealth information. The lab contains three tasks: Visual-Interactive Search and Exploration of eHealth Data, Information extraction from clinical text, User-centered health information retrieval.

ImageCLEF: ImageCLEF aims at providing benchmarks for the challenging task of image annotation for a wide range of source images and annotation objective. The tasks address different aspects of the annotation problem and are aimed at supporting and promoting the cutting-edge research addressing the key challenges in the field, such as multi-modal image annotation, domain adaptation and ontology driven image annotation. The Lab tasks are: Robot Vision, Scalable concept Image Annotation, Liver CT Annotation and Domain Adaptation.

INEX — Initiative for the Evaluation of XML retrieval: INEX builds evaluation benchmarks for search in the context of rich structure such as document structure, semantic metadata, entities, or genre/topical structure. INEX 2014 runs the following tasks: Social Book Search Task, Interactive Social Book Search Task, Tweet Contextualization Task.

LifeCLEF: LifeCLEF aims at evaluating multimedia analysis and retrieval techniques on biodiversity data for species identification. The three Tasks regard respectively birds, plants and fishes: BirdCLEF (a bird songs identification task based on Xeno-Canto audio recordings), PlantCLEF (an image-based plant identification task based on the data of Tela Botanica social network) and FishCLEF (a fish video surveillance task based on the data of the Fish4Knowledge network).

NEWSREEL - News Recommendation Evaluation Lab: Newsreel Lab is the first news recommendation evaluation Lab and it is divided in two different tasks. One task is about predicting the items a user will click in the next 10 Minutes based on the offline dataset. The other task is about predicting the articles users will click; in this task, the prediction algorithms are evaluated in an online scenario based on live user-interactions.

PAN Lab on Uncovering Plagiarism, Authorship, and Social Software Misuse: PAN centers around the topics of plagiarism, authorship, and social software misuse. The goal is to foster research on automatic detection and uncovering. People increasingly share their work online, contribute to open projects and engage in web-based social interactions. The ease and anonymity with which this can be done raises concerns about verifiability and trust: Is a given text an original? Is the author the one

who she claims to be? Does a piece of information come from a trusted source? There are three Lab Tasks: Author Identification, Author Profiling and Plagiarism Detection.

QA Track — CLEF Question Answering Track: In the current general scenario for the CLEF QA Track, the starting point is always a Natural Language question. However, answering some questions may need to query Linked Data; whereas some questions may need textual inferences and querying free-text. Answering some queries may need both. The tasks are: QALD - Question Answering over Linked Data, Bio-ASQ - Biomedical semantic indexing and question answering, Entrance Exams.

RepLab: The aim of RepLab is to bring together the Information Access research community with representatives from the Online Reputation Management industry, with the ultimate goals of (i) establishing a roadmap on the topic that includes a description of the language technologies required in terms of resources, algorithms, and applications; (ii) specifying suitable evaluation methodologies and metrics to measure scientific progress; and (iii) developing of test collections that enable systematic comparison of algorithms and reliable benchmarking of commercial systems. Lab Tasks are Reputation Dimensions and Author Profiling.

Acknowledgements

We would like to thank the members of CLEF-LOC (the CLEF Lab Organization Committee) for their thoughtful and elaborate contributions to assessing the proposals during the selection process:

Donna Harman, National Institute for Standards and Technology (NIST), USA

Carol Peters, ISTI, National Council of Research (CNR), Italy

Marteen de Rijke, University of Amsterdam, The Netherlands

Last but not least without the important and tireless effort of the enthusiastic and creative proposal authors, the organizers of the selected labs, the colleagues and friends involved in running them, and the participants who contribute their time to making the labs and workshops a success, the CLEF labs would not be possible.

Thank you all very much!

July 2014

Martin Halvey and Wessel Kraaij
Lab Chairs

Linda Cappellato and Nicola Ferro
Working Notes Editors