

Preface

The CLEF 2015 conference is the sixteenth 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.

Eight laboratories were selected and run during CLEF 2015. 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 Vienna on March 31st 2015 and supported by the ELIAS network¹.

CLEF has been always backed by European projects which complement the incredible amount of volunteering work performed by Lab Organizers and the CLEF community with the resources needed for its necessary central coordination, in a similar manner to the other major international evaluation initiatives as TREC, NTCIR, FIRE and MediaEval. Since 2014, the organisation of CLEF no longer has direct support from European projects and working to transform itself into a self-sustainable activity. This is being made 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.

The Labs at CLEF 2015, building on previous experience, demonstrate the maturity of the CLEF evaluation environment via the incorporation of new tasks, new and larger data sets, new ways of evaluation or more languages. Details of the individual Labs are described by the Lab organizers in these proceedings, here we just provide brief comment on each one.

¹ <http://www.elias-network.eu/>

CLEFeHealth: CLEFeHealth explores scenarios which aim to ease patients and nurses understanding and accessing of eHealth information. The goals of the lab are to develop processing methods and resources in a multilingual setting to enrich difficult-to-understand eHealth texts, and provide valuable documentation. The lab contains two tasks: Information Extraction from Clinical Data (with two subtasks: Clinical speech recognition and Named entity recognition from clinical narratives in European languages) and User-centered Health Information Retrieval (with two subtasks: Monolingual IR and Multilingual IR).

ImageCLEF: In 2015, ImageCLEF organized four main tasks (Image Annotation, Medical Classification, Medical Clustering, Liver CT Annotation) with a global objective of benchmarking automatic annotation and indexing of images. The tasks tackle different aspects of the annotation problem and are aimed at supporting and promoting cutting-edge research addressing the key challenges in the field:

Life CLEF: The LifeCLEF lab continues image-based plant identification task which was originally run within ImageCLEF since 2011, with the same tasks of last year (BirdCLEF, PlantCLEF and FishCLEF). However, the LifeCLEF tasks radically enlarges the evaluated challenge towards multimodal data by (i) considering birds and fish in addition to plants, (ii) considering audio and video content in addition to images, (iii) scaling-up the evaluation data to hundreds of thousands of life media records and thousands of living species.

LL4IR - Living Labs for IR: CLEF 2015 sees the first edition of this new lab, which features one task (Product search and web search). The main goal of the Lab is to provide a benchmarking platform for researchers to evaluate their ranking systems in a live setting with real users in their natural task environments. The lab acts as a proxy between commercial organizations (live environments) and lab participants (experimental systems), facilitates data exchange, and makes comparison between the participating systems.

NEWSREEL – News Recommendation Evaluation Lab: CLEF 2015 is the second iteration of this lab. Participants can: a) develop news recommendation algorithms and b) have them tested by millions of users over the period of a few weeks in a living lab. NEWSREEL provides two tasks designed to address the challenge of real-time news recommendation: Benchmark News Recommendations in a Living Lab and Benchmarking News Recommendations in a Simulated Environment.

PAN – Uncovering Plagiarism, Authorship and Social Software Misuse: This is the 12th edition of the PAN lab on evaluation of uncovering plagiarism, authorship, and social software misuse. PAN offers three tasks at CLEF 2015 with new evaluation resources consisting of large-scale corpora, performance measures, and web services that allow for meaningful evaluations. The main goal is to provide for sustainable and reproducible evaluations, to get a clear view of the capabilities of state-of-the-art-algorithms. The tasks are: Plagiarism Detection, Author Identification, Author Profiling.

QA – Question answering: 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 (especially if aggregations or logical inferences are required); 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), Entrance Exams (Questions from reading tests), BioASQ (Large-Scale Biomedical Semantic Indexing), and BioASQ (Biomedical Question answering).

SBS – Social Book Search: The Social Book Search Lab was previously part of the INEX evaluation benchmark (since 2007). Real-world information needs are generally complex, yet almost all research focuses instead on either relatively simple search based on queries or recommendation based on profiles. The goal of the Social Book Search Lab is to investigate techniques to support users in complex book search tasks that involve more than just a query and results list. SBS runs two tasks: Suggestion Track and Interactive Track.

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

Jacques Savoy, University of Neuchâtel, Switzerland

William Webber, William Webber Consulting, Australia

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 2015

Gareth J. F. Jones and Eric San Juan
Lab Chairs

Linda Cappellato and Nicola Ferro
Working Notes Editors