

What happened in CLEF 2007?

Introduction to the Working Notes

Carol Peters

Istituto di Scienza e Tecnologie dell'Informazione (ISTI-CNR), Pisa, Italy
carol.peters@isti.cnr.it

The objective of the Cross Language Evaluation Forum¹ is to promote research in the field of multilingual system development. This is done through the organisation of annual evaluation campaigns in which a series of tracks designed to test different aspects of mono- and cross-language information retrieval (IR) are offered. The intention is to encourage experimentation with all kinds of multilingual information access – from the development of systems for monolingual retrieval operating on many languages to the implementation of complete multilingual multimedia search services. This has been achieved by offering an increasingly complex and varied set of evaluation tasks over the years. The aim is not only to meet but also to anticipate the emerging needs of the R&D community and to encourage the development of next generation multilingual IR systems.

These Working Notes contain descriptions of the experiments conducted within CLEF 2007 – the eighth in a series of annual system evaluation campaigns. The results of the experiments will be presented and discussed in the CLEF 2007 Workshop, 19-21 September, Budapest, Hungary. The final papers - revised and extended as a result of the discussions at the Workshop - together with a comparative analysis of the results will appear in the CLEF 2007 Proceedings, to be published by Springer in their Lecture Notes for Computer Science series.

As from CLEF 2005, the Working Notes are published in electronic format only and are distributed to participants at the Workshop on CD-ROM together with the Book of Abstracts in printed form. All reports included in the Working Notes will also be inserted in the DELOS Digital Library, accessible at <http://delos-dl.isti.cnr.it>.

Both Working Notes and Book of Abstracts are divided into eight sections, corresponding to the CLEF 2007 evaluation tracks, plus an additional section describing other evaluation initiatives using CLEF data: MorphoChallenge 2007 and SemEval 2007. In addition appendices are included containing run statistics for the Ad Hoc, Domain-Specific, GeoCLEF and CL-SR tracks, plus a list of all participating groups showing in which track they took part.

The main features of the 2007 campaign are briefly outlined here below in order to provide the necessary background to the experiments reported in the rest of the Working Notes.

1. Tracks and Tasks in CLEF 2007

CLEF 2007 offered seven tracks designed to evaluate the performance of systems for:

- mono-, bi- and multilingual textual document retrieval on news collections (Ad Hoc)
- mono- and cross-language information on structured scientific data (Domain-Specific)
- multiple language question answering (QA@CLEF)
- cross-language retrieval in image collections (ImageCLEF)
- cross-language speech retrieval (CL-SR)
- multilingual retrieval of Web documents (WebCLEF)
- cross-language geographical retrieval (GeoCLEF)

These tracks are mainly the same as those offered in CLEF2006 with the exclusion of an interactive track², however many of the tasks offered are new.

¹ CLEF is included in the activities of the DELOS Network of Excellence on Digital Libraries, funded by the Sixth Framework Programme of the European Commission. For information on DELOS, see www.delos.info.

² From CLEF 2001 through CLEF 2006, we have offered an interactive track. Unfortunately, this year, the track was suspended due to other commitments of the organisers. Owing to the importance of user intervention in cross-language IR, we intend to re-propose and strengthen the interactive activity in CLEF 2008.

Cross-Language Text Retrieval (Ad Hoc): This year, this track offered **mono-** and **bilingual** tasks on target collections for central European languages (Bulgarian, Czech³ and Hungarian). Similarly to last year, a bilingual task encouraging system testing with non-European languages against English documents was offered. Topics were made available in Amharic, Chinese, Oromo and Indonesian. A special sub-task regarded Indian language search against an English target collection was also organised with the assistance of a number of Indian research institutes, responsible for the preparation of the topics. The languages offered were Hindi, Bengali, Tamil, Telugu and Marathi. In order to establish benchmarks in this subtask, all participating groups has to submit:

- one monolingual English to English run (mandatory)
- at least one run in Hindi to English (mandatory)
- runs in other Indian languages to English (optional).

A "robust" task was again be offered, emphasizing the importance of reaching a minimal performance for all topics instead of high average performance. Robustness is a key issue for the transfer of CLEF research into applications. The 2007 robust task involved three languages often used in previous CLEF campaigns (English, French, Portuguese). The track was coordinated jointly by ISTI-CNR and U.Padua (Italy) and U.Hildesheim (Germany).

Cross-Language Scientific Data Retrieval (Domain-Specific): Mono- and cross-language domain-specific retrieval was studied in the domain of social sciences using structured data (e.g. bibliographic data, keywords, and abstracts) from scientific reference databases. The target collections provided were: GIRT-4 for German/English, INION for Russian and Cambridge Sociological Abstracts for English. A multi-lingual controlled vocabulary (German, English, Russian) suitable for use with GIRT-4 and INION together with a bi-directional mapping between this vocabulary and that used for indexing the Sociological Abstracts (English) was provided. Topics were offered in English, German and Russian. This track was coordinated by IZ Bonn (Germany).

Multilingual Question Answering (QA@CLEF): QA@CLEF 2007 proposed both main and pilot tasks. The main task scenario was topic-related QA, where the questions are grouped by topics and may contain anaphoric references one to the others. The answers were retrieved from heterogeneous document collections, i.e. news articles and Wikipedia. Many sub-tasks were set up, monolingual – where the questions and the target collections searched for answers are in the same language - and bilingual – where source and target languages are different. Bulgarian, Dutch, English, French, German, Italian, Portuguese, Romanian and Spanish were offered as target languages; query languages used in the bilingual tasks depended on demand (see the track overview for details). Following the positive response at QA@CLEF 2006, the Answer Validation Exercise (AVE) was repropose. A new pilot tasks was also offered: Question Answering on Speech Transcript (QAst), in which the answers to factual questions have to be extracted from spontaneous speech transcriptions (manual and automatic transcriptions) coming from different human interaction scenario. The track is organized by several institutions (one for each source language) and jointly coordinated by CELCT, Trento (Italy), LSI-UNED, Madrid and UPC, Barcelona (Spain).

Cross-Language Retrieval in Image Collections (ImageCLEF): This track evaluated retrieval of images described by text captions in several languages; both text and image retrieval techniques were exploitable. Four challenging tasks were offered: (i) multilingual ad-hoc retrieval (collection with mixed English/German/Spanish annotations, queries in more languages), (ii) medical image retrieval (casenotes in English/ French/German; visual, mixed, semantic queries in same languages), (iii) hierarchical automatic image annotation for medical images (fully categorized in English and German, purely visual task), (iv) photographic annotation through detection of objects in images (using the same collection as (i) with a restricted number of objects, a purely visual task). Image retrieval was not required for all tasks and a default visual and textual retrieval system was made available for participants. The track coordinators were U.Sheffield (UK) and the U. and U. Hospitals of Geneva (Switzerland). Oregon Health and Science U. (US), Victoria U., Melbourne (Australia), RWTH Aachen (Germany) and Vienna Univ. Tech (Austria) collaborated in the task organization.

Cross-Language Speech Retrieval (CL-SR): The focus is on searching spontaneous speech from oral history interviews rather than news broadcasts. The test collection created for the track is a subset of a large archive of videotaped oral histories from survivors, liberators, rescuers and witnesses of the Holocaust created by the Survivors of the Shoah Visual History Foundation (VHF). Automatic Speech Recognition (ASR) transcripts and both automatically assigned and manually assigned thesaurus terms were available as part of the collection.

In 2006 the CL-SR track included search collections of conversational English and Czech speech using six languages (Czech, Dutch, English, French, German and Spanish). In CLEF 2007 additional topics were added for the Czech speech collection. Speech content is described by automatic speech transcriptions manually and automatically assigned controlled vocabulary descriptors for concepts, dates and locations, manually assigned

³ New this year.

person names, and hand-written segment summaries. The track was coordinated by U. Maryland (USA), Dublin City U. (Ireland) and Charles U. (Czech Republic).

Multilingual Web Retrieval (WebCLEF): The WebCLEF 2007 task combines insights gained from previous editions of WebCLEF 2005–2006 and the WiQA 2006 pilot, and goes beyond the navigational queries considered at WebCLEF 2005 and 2006. At WebCLEF 2007 so-called undirected informational search goals were considered in a web setting: “I want to learn anything/everything about my topic.” The track was coordinated by U. Amsterdam (The Netherlands).

Cross-Language Geographical Retrieval (GeoCLEF): The purpose of GeoCLEF is to test and evaluate cross-language geographic information retrieval (GIR): retrieval for topics with a geographic specification. GeoCLEF 2007 consisted of two sub tasks. A search task ran for the third time and a query classification task was organized for the first. For the GeoCLEF 2007 search task, twenty-five search topics were defined by the organizing groups for searching English, German, Portuguese and Spanish document collections. Topics were translated into English, German and Spanish. For the classification task, a query log from a search engine was provided and the groups needed to identify the queries with a geographic scope and the geographic components within the local queries. The track was coordinated jointly by UC Berkeley (USA), U.Sheffield (UK), U. Hildesheim (Germany), Linguatca SINTEF (Norway), Microsoft Asia (China).

Details on the technical infrastructure and the organisation of these tracks can be found in the track overview reports in this volume, collocated at the beginning of the relevant sections.

2. Test Collections

A number of different document collections were used in CLEF 2007 to build the test collections:

- CLEF multilingual comparable corpus of more than 3 million news documents in 13 languages; new data was added this year for Czech, Bulgarian and English (see Table 1); Parts of this collections were used in the Ad-Hoc, QuestionAnswering, and GeoCLEF tracks.
- The GIRT-4 social science database in English and German (over 300,000 documents) and two Russian databases: the Russian Social Science Corpus (approx. 95,000 documents) and the Russian ISISS collection for sociology and economics (approx. 150,000 docs). The RSSC corpus was not used this year. Cambridge Sociological Abstracts in English. These collections were used in the domain-specific track.
- The ImageCLEF track used collections for both general photographic and medical image retrieval:
 - IAPR TC-12 photo database of 25,000 photographs with captions in English, German and Spanish; PASCAL VOC 2006 training data (new this year);
 - ImageCLEFmed radiological database consisting of 6 distinct datasets – 2 more than last year; IRMA collection in English and German of 12,000 classified images for automatic medical image annotation
- Malach collection of spontaneous conversational speech derived from the Shoah archives in English (more than 750 hours) and Czech (approx 500 hours). This collection was used in the speech retrieval track.
- EuroGOV, a multilingual collection of about 3.5M webpages, containing documents many languages crawled from European governmental sites, used in the WebCLEF track.

3. Technical Infrastructure

The CLEF technical infrastructure is managed by the DIRECT system. DIRECT manages the test data plus results submission and analyses for the ad hoc, question answering and geographic IR tracks. It has been designed to facilitate data management tasks but also to support the production, maintenance, enrichment and interpretation of the scientific data for subsequent in-depth evaluation studies.

The technical infrastructure is thus responsible for:

- the track set-up, harvesting of documents, management of the registration of participants to tracks;
- the submission of experiments, collection of metadata about experiments, and their validation;
- the creation of document pools and the management of relevance assessment;
- the provision of common statistical analysis tools for both organizers and participants in order to allow the comparison of the experiments;
- the provision of common tools for summarizing, producing reports and graphs on the measured performances and conducted analyses.

DIRECT is designed and implemented by Giorgio Di Nunzio and Nicola Ferro

Table 1: Sources and dimensions of the CLEF 2007 multilingual comparable corpus

Collection	Added in	Size (MB)	No. of Docs	Median Size of Docs. (Bytes)	Median Size of Docs. (Tokens) ⁴	Median Size of Docs (Features)
Bulgarian: Sega 2002	2005	120	33,356	NA	NA	NA
Bulgarian: Standart 2002	2005	93	35,839	NA	NA	NA
Bulgarian: Novinar 2002	2007	48	18,086	NA	NA	NA
Czech: Mladna frontaDnes 2002	2007	143	68,842	NA	NA	NA
Czech: Lidove Noviny 2002	2007	35	12,893	NA	NA	NA
Dutch: Algemeen Dagblad 94/95	2001	241	106483	1282	166	112
Dutch: NRC Handelsblad 94/95	2001	299	84121	2153	354	203
English: LA Times 94	2000	425	113005	2204	421	246
English: LA Times 2002	2007	434	135,153	NA	NA	NA
English: Glasgow Herald 95	2003	154	56472	2219	343	202
Finnish: Aamulehti late 94/95	2002	137	55344	1712	217	150
French: Le Monde 94	2000	158	44013	1994	361	213
French: ATS 94	2001	86	43178	1683	227	137
French: ATS 95	2003	88	42615	1715	234	140
German: Frankfurter Rundschau94	2000	320	139715	1598	225	161
German: Der Spiegel 94/95	2000	63	13979	1324	213	160
German: SDA 94	2001	144	71677	1672	186	131
German: SDA 95	2003	144	69438	1693	188	132
Hungarian: Magyar Hirlap 2002	2005	105	49,530	NA	NA	NA
Italian: La Stampa 94	2000	193	58051	1915	435	268
Italian: AGZ 94	2001	86	50527	1454	187	129
Italian: AGZ 95	2003	85	48980	1474	192	132
Portuguese: Público 1994	2004	164	51751	NA	NA	NA
Portuguese: Público 1995	2004	176	55070	NA	NA	NA
Portuguese: Folha 94	2005	108	51,875	NA	NA	NA
Portuguese: Folha 95	2005	116	52,038	NA	NA	NA
Russian: Izvestia 95	2003	68	16761	NA	NA	NA
Spanish: EFE 94	2001	511	215738	2172	290	171
Spanish: EFE 95	2003	577	238307	2221	299	175
Swedish: TT 94/95	2002	352	142819	2171	183	121

SDA/ATS/AGZ = Schweizerische Depeschagentur (Swiss News Agency)

EFE = Agencia EFE S.A (Spanish News Agency)

TT = Tidningarnas Telegrambyrå (Swedish newspaper)

⁴ The number of tokens extracted from each document can vary slightly across systems, depending on the respective definition of what constitutes a token. Consequently, the number of tokens and features given in this table are approximations and may differ from actual implemented systems.

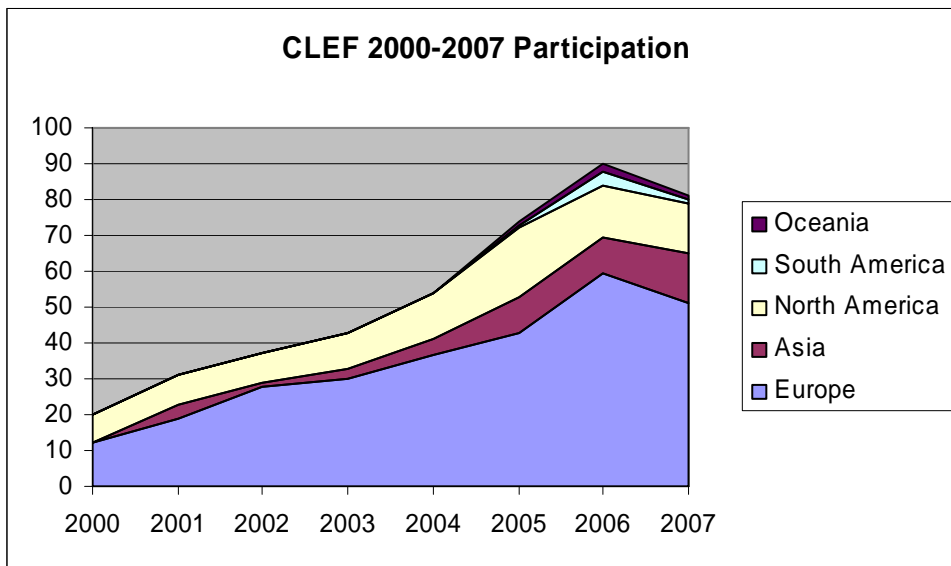


Figure 1. CLEF 2000 – 2007: Variation in Participation




Figure 2. CLEF 2000 – 2007: Participation per Track in Tracks

4. Participation

A total of 81 groups submitted runs in CLEF 2006, slightly down from the 90 groups of CLEF 2005: 51(59.5) from Europe, 14(14.5) from N.America; 14(10) from Asia, 1(4) from S.America and 1(1) from Australia. The breakdown of participation of groups per track is as follows: Ad Hoc 22(25); Domain-Specific 5(4); QAatCLEF 28(37); ImageCLEF 35(25); CL-SR 8(6); WebCLEF 4(8); GeoCLEF 13(17)⁵. A list of groups and indications of the tracks in which they participated is given in the Appendix to these Working Notes. Figure 1 shows the variation in participation over the years and Figure 2 shows the shift in focus as new tracks have been added

⁵ Last year's figures are between brackets.

In particular, these figures show that while there is a constant increase in interest in the ImageCLEF track, there is a consistent decrease in popularity of the question answering and web tracks. Although the fluctuation in QA does not seem to be of great significance - this is a very difficult task - the apparent lack of interest in WebCLEF is surprising. With the importance of Internet and web search engines, a larger participation in this task is to be expected. The large numbers for ImageCLEF also give rise to some discussion. The defining feature of CLEF is its multilinguality; ImageCLEF is perhaps the least multilingual of the CLEF tracks as much of the work is done in a language-independent context. These questions will be the subject of debate at the workshop. At the same time, it should be noted that these Working Notes also include reports from two separate evaluation initiatives which actually used CLEF data for certain tasks – thus the impact of CLEF spreads far beyond the boundaries of the CLEF evaluation campaigns.

5. Workshop

CLEF aims at creating a strong CLIR/MLIR research and development community. The Workshop plays an important role by providing the opportunity for all the groups that have participated in the evaluation campaign to get together comparing approaches and exchanging ideas. The work of the groups participating in this year's campaign will be presented in plenary paper and poster sessions. There will also be break-out sessions for more in-depth discussion of the results of individual tracks and intentions for the future. The final sessions will include discussions on ideas for new tracks in future campaigns. Overall, the Workshop should provide an ample panorama of the current state-of-the-art and the latest research directions in the multilingual information retrieval area. I very much hope that it will prove an interesting, worthwhile and enjoyable experience to all those who participate.

The final programme and the presentations at the Workshop will be posted on the CLEF website at <http://www.clef-campaign.org>.

Acknowledgements

It would be impossible to run the CLEF evaluation initiative and organize the annual workshops without considerable assistance from many groups. CLEF is organized on a distributed basis, with different research groups being responsible for the running of the various tracks. My gratitude goes to all those who have been involved in the coordination of the 2007 campaigns. A list of the main institutions involved is given on the following page. Here below, let me thank the people mainly responsible for the coordination of the different tracks:

- Giorgio Di Nunzio, Nicola Ferro and Thomas Mandl for the Ad Hoc Track
- Vivien Petras, Stefan Baerisch, Maximillian Stempfhuber for the Domain-Specific track
- Bernardo Magnini, Danilo Giampiccolo, Pamela Forner, Anselmo Peñas, Christelle Ayache, Corina Forăscu, Valentin Jijkoun, Petya Osenova, Paulo Rocha, Bogdan Sacaleanu, Richard Sutcliffe for QA@CLEF
- Allan Hanbury, Paul Clough, Henning Müller, Thomas Deselaers, Michael Grubinger, Jayashree Kalpathy-Cramer and William Hersh for ImageCLEF
- Douglas W. Oard, Gareth J. F. Jones, and Pavel Pecina for CL-SR
- Valentin Jijkoun and Maarten de Rijke for Web-CLEF
- Thomas Mandl, Fredric Gey, Giorgio Di Nunzio, Nicola Ferro, Ray Larson, Mark Sanderson, Diana Santos, Christa Womser-Hacker, Xing Xie for GeoCLEF

I also thank all those colleagues who have helped us by preparing topic sets in different languages and in particular the NLP Lab. Dept. of Computer Science and Information Engineering of the National Taiwan University for their work on Chinese.

I should also like to thank the members of the CLEF Steering Committee who have assisted me with their advice and suggestions throughout this campaign.

Furthermore, I gratefully acknowledge the support of all the data providers and copyright holders, and in particular:

- The Los Angeles Times, for the American-English data collection
- SMG Newspapers (The Herald) for the British-English data collection
- Le Monde S.A. and ELDA: Evaluations and Language resources Distribution Agency, for the French data
- Frankfurter Rundschau, Druck und Verlagshaus Frankfurt am Main; Der Spiegel, Spiegel Verlag, Hamburg, for the German newspaper collections
- InformationsZentrum Sozialwissenschaften, Bonn, for the GIRT database
- SocioNet system for the Russian Social Science Corpora
- Hypersystems Srl, Torino and La Stampa, for the Italian data

- Agencia EFE S.A. for the Spanish data
- NRC Handelsblad, Algemeen Dagblad and PCM Landelijke dagbladen/Het Parool for the Dutch newspaper data
- Aamulehti Oyj and Sanoma Osakeyhtiö for the Finnish newspaper data
- Russika-Izvestia for the Russian newspaper data
- Público, Portugal, and Linguateca for the Portuguese (PT) newspaper collection
- Folha, Brazil, and Linguateca for the Portuguese (BR) newspaper collection
- Tidningarnas Telegrambyrå (TT) SE-105 12 Stockholm, Sweden for the Swedish newspaper data
- Schweizerische Depeschagentur, Switzerland, for the French, German and Italian Swiss news agency data
- Ringier Kiadoi Rt. [Ringier Publishing Inc.] and the Research Institute for Linguistics, Hungarian Acad. Sci. for the Hungarian newspaper documents
- Sega AD, Sofia; Standart Nyuz AD, Novinar OD Sofia, and the BulTreeBank Project, Linguistic Modelling Laboratory, IPP, Bulgarian Acad. Sci, for the Bulgarian newspaper documents
- Mafra a.s. and Lidové Noviny a.s. for the Czech newspaper data
- St Andrews University Library for the historic photographic archive
- University and University Hospitals, Geneva, Switzerland and Oregon Health and Science University for the ImageCLEFmed Radiological Medical Database
- The Radiology Dept. of the University Hospitals of Geneva for the Casimage database and the PEIR (Pathology Education Image Resource) for the images and the HEAL (Health Education Assets Library) for the Annotation of the Peir dataset.
- Aachen University of Technology (RWTH), Germany for the IRMA database of annotated medical images
- Mallinkrodt Institute of Radiology for permission to use their nuclear medicine teaching file
- University of Basel's Pathopic project for their Pathology teaching file
- Michael Grubinger, administrator of the IAPR Image Benchmark, Clement Leung who initiated and supervised the IAPR Image Benchmark Project, and André Kiwitz, the Managing Director of Viventura for granting access to the image database and the raw image annotations of the tour guides.
- The Survivors of the Shoah Visual History Foundation, and IBM for the Malach spoken document collection

Without their contribution, this evaluation activity would be impossible.

Last and not least, I should like to express my gratitude to Alessandro Nardi and Valeria Quochi for their assistance in the organisation of the CLEF 2007 Workshop.

Coordination

CLEF is coordinated by the Istituto di Scienza e Tecnologie dell'Informazione, Consiglio Nazionale delle Ricerche, Pisa. The following institutions have contributed to the organisation of the different tracks of the CLEF 2007 campaign:

- Centre for the Evaluation of Human Language and Multimodal Communication Technologies (CELCT), Trento, Italy
- College of Information Studies and Institute for Advanced Computer Studies, University of Maryland, USA
- Department of Computer Science, University of Indonesia
- Department of Computer Science Department, RWTH Aachen University, Germany
- Department of Computer Science and Information Systems, University of Limerick, Ireland
- Department of Computer Science and Information Engineering, National University of Taiwan
- Department of Information Engineering, University of Padua, Italy
- Department of Information Science, University of Hildesheim, Germany
- Department of Information Studies, University of Sheffield, UK
- Department of Medical Informatics, Aachen University of Technology (RWTH), Germany
- Evaluations and Language Resources Distribution Agency Sarl, Paris, France
- Fondazione Bruno Kessler FBK-irst, Trento, Italy
- German Research Centre for Artificial Intelligence, DFKI, Saarbrücken,
- Information and Language Processing Systems, University of Amsterdam, Netherlands
- InformationsZentrum Sozialwissenschaften, Bonn, Germany
- Institute for Information Technology, Hyderabad, India
- Institute of Formal and Applied Linguistics, Charles University, Czech Rep
- Lenguajes y Sistemas Informáticos, Universidad Nacional de Educación a Distancia, Madrid, Spain
- Linguateca, Sintef, Oslo, Norway; University of Minho, Braga, Portugal
- Linguistic Modelling Laboratory, Bulgarian Academy of Sciences
- Microsoft Research Asia
- National Institute of Standards and Technology, Gaithersburg MD, USA
- Oregon Health and Science University, USA
- Research Computing Center of Moscow State University
- Research Institute for Linguistics, Hungarian Academy of Sciences
- School of Computer Science and Mathematics, Victoria University, Australia
- School of Computing, Dublin City University, Ireland
- UC Data Archive and School of Information Management and Systems, UC Berkeley, USA
- University "Alexandru Ioan Cuza", IASI, Romania
- University Hospitals and University of Geneva, Switzerland
- Vienna University of Technology, Austria

CLEF Steering Committee

Maristella Agosti, University of Padova, Italy
Martin Braschler, Zurich University of Applied Sciences Winterthur, Switzerland
Amedeo Cappelli, ISTI-CNR & CELCT, Italy
Hsin-Hsi Chen, National Taiwan University, Taipei, Taiwan
Khalid Choukri, Evaluations and Language resources Distribution Agency, Paris, France
Paul Clough, University of Sheffield, UK
Thomas Deselaers, RWTH Aachen University, Germany
David A. Evans, Clairvoyance Corporation, USA
Marcello Federico, ITC-irst, Trento, Italy
Christian Fluhr, CEA-LIST, Fontenay-aux-Roses, France
Norbert Fuhr, University of Duisburg, Germany
Frederic C. Gey, U.C. Berkeley, USA
Julio Gonzalo, LSI-UNED, Madrid, Spain
Donna Harman, National Institute of Standards and Technology, USA
Gareth Jones, Dublin City University, Ireland
Franciska de Jong, University of Twente, Netherlands
Noriko Kando, National Institute of Informatics, Tokyo, Japan
Jussi Karlgren, Swedish Institute of Computer Science, Sweden
Michael Kluck, German Institute for International and Security Affairs, Berlin, Germany
Natalia Loukachevitch, Moscow State University, Russia
Bernardo Magnini, ITC-irst, Trento, Italy
Paul McNamee, Johns Hopkins University, USA
Henning Müller, University & University Hospitals of Geneva, Switzerland
Douglas W. Oard, University of Maryland, USA
Maarten de Rijke, University of Amsterdam, Netherlands
Diana Santos, Linguatca, Sintef, Oslo, Norway
Jacques Savoy, University of Neuchatel, Switzerland
Peter Schäuble, Eurospider Information Technologies, Switzerland
Richard Sutcliffe, University of Limerick, Ireland
Max Stempfhuber, Informationszentrum Sozialwissenschaften Bonn, Germany
Hans Uszkoreit, German Research Center for Artificial Intelligence (DFKI), Germany
Felisa Verdejo, LSI-UNED, Madrid, Spain
José Luis Vicedo, University of Alicante, Spain
Ellen Voorhees, National Institute of Standards and Technology, USA
Christa Womser-Hacker, University of Hildesheim, Germany