



**Project no. 215231**

**TrebleCLEF**

Treble-CLEF: Evaluation, Best Practices and Collaboration for Multilingual Information Access  
IST: ICT-1-4-1, Digital libraries and technology-enhanced learning

**Deliverable 3.1**

**Bringing Multilingual Information Access to Operational Systems  
– TrebleCLEF System Developers Workshop  
REPORT**

Start Date of Project: 01 January 2008

Duration: 24 Months

Organisation Name of Lead Contractor for this Deliverable: ZHAW

Version 1.00, December 2008 [Final]

Project co-funded by the European Commission within the Seventh Framework programme

---

## Document Information

Deliverable number: 3.1  
Deliverable title: Bringing Multilingual Information Access to Operational System – TrebleCLEF System Developers Workshop REPORT  
Due date of deliverable: 31/12/2008  
Actual date of deliverable: 22/12/2008  
Author(s): Martin Braschler, Paul Clough  
Participant(s): ZHAW, USFD  
Workpackage: 3  
Workpackage title: Best Practices in System Development and User Studies  
Workpackage leader: UNED  
Dissemination Level: Public  
Version: 1.00 final  
Keywords: System Developer Workshop, Operational Multilingual Information Access, Technology Transfer

### History of Versions

Version	Date	Status	Author (Partner)	Description/Approval Level
1.00	22/12/08	final	ZHAW, USFD	

## Abstract

This document summarises the System Developers Workshop entitled “Bringing Multilingual Information Access to Operational Systems” held 2-3 October 2008 at Zurich University of Applied Sciences (ZHAW) in Winterthur, Switzerland. A number of prominent representatives from both academic and commercial institutions participated in the event and discussed the more operational/applied aspects of Multilingual Information Access, CLEF and test collections for system evaluation. This document summarises the workshop audience, goals, participant presentations and main discussion points generated from the discussions. Together with the actual workshop, it forms deliverable 3.1 of the TrebleCLEF project.

## Table of Contents

Document Information .....	1
Abstract .....	1
Executive Summary .....	3
1 Introduction .....	4
2 Audience.....	4
3 Goals.....	5
4 Presentations.....	5
5 Discussions.....	8
5.1 Questions .....	8
5.2 Category CLEF tasks .....	10
5.3 Resources for MLIA .....	11
5.4 Components for MLIA .....	11
5.5 Varia.....	11
6 Summary .....	11
Acknowledgements .....	12
References .....	12
Appendix – Attendees (alphabetic listing) .....	13

## Executive Summary

This document summarises the System Developers Workshop entitled “Bringing Multilingual Information Access to Operational Systems” held 2-3 October 2008 at Zurich University of Applied Sciences (ZHAW) in Winterthur, Switzerland. A number of prominent representatives from both academic and commercial institutions participated in the event and discussed the more operational/applied aspects of Multilingual Information Access, CLEF and test collections for system evaluation. Despite continued interest in multilingualism within the Digital Library community (and beyond); there still exist few operational systems. The aim of this workshop was to try and understand the reasons behind this apparent gap, in particular from the system developer’s perspective.

All attendees provided input to the event through presentations and in subsequent discussions. This will be used as source material for future analysis by the TrebleCLEF consortium. Although the workshop included participants with varying academic and commercial interests, all attendees shared a background of (sometimes limited) involvement with CLEF (either as organisers or participants) and could therefore provide helpful insights both from an external and internal perspective on the use of CLEF results within an operational context. Issues raised included the realism of the evaluation resources with respect to use cases and scenarios, performance measures, verifiability and robustness of results, presentation of multilingual results and user-orientated evaluation, and the transferability of results to operational settings.

Presentations at the workshop provided a useful stimulus for further discussions, the result of which was a list of questions regarding CLEF and the evaluation of multilingual information access systems. These have been summarised in this deliverable, together with an initial discussion of their implications for future CLEF events. The output from the System Developers Workshop is one of the key sources of information which will be used to compile best practice recommendations for system-orientated evaluation. These best-practice recommendations, based on a thorough analysis of the results of the workshop, will be contained in the forthcoming deliverable 3.3.

## 1 Introduction

The TrebleCLEF System Developers Workshop, titled "Bringing Multilingual Information Access to Operational Systems" was held on October 2 and 3, 2008 in Winterthur near Zurich, Switzerland [3]. The event was hosted by the Zurich University of Applied Sciences (ZHAW). This workshop, together with a summary of workshop activities, constitutes deliverable 3.1 of the TrebleCLEF coordination action [1].

The global availability of networked information sources and digital libraries have led to a strong demand for multilingual access and communication technologies. There have been notable recent research advances, evidenced for example in the work carried out by the many participants of the CLEF evaluation campaigns from 2000-2008 [2]. However, despite these advances, there are still few operational systems in existence, and these are usually limited to the most widely used languages.

The System Developers Workshop was thus designed to cover questions of encouraging transfer of research results and disseminating know-how, tools, resources and best practice guidelines in the field of Multilingual Information Access (MLIA) and Cross-Language Information Retrieval (CLIR). The output from the workshop is one of the key sources of information for the later compilation of best practice recommendations for system-oriented MLIA by the TrebleCLEF consortium (Deliverable 3.3).

The remainder of this deliverable is structured as follows: the setup of the workshop (audience and goals; sections 2 and 3) is presented first, followed by a summary of presentations given by attendees at the workshop (section 4). A focus of the workshop was to stimulate intensive discussion among attendees regarding open questions related to MLIA. In section 5, we first present an extensive list of such open questions as identified by the attendees, grouped into four different subject areas. We then proceed to give a summary of discussions from the final plenary session for each of these subject areas. Section 6 concludes the deliverable by giving a summary and outlook on Deliverable 3.3, which will include an analysis of the output from this workshop.

The present document assumes a familiarity by the reader of MLIA and information retrieval terminology.

## 2 Audience

Attendance at the workshop was by invitation only and attendees were selected to be from one of the following groups:

- Current and past participants of the CLEF campaigns that either have a commercial background in MLIA/CLIR or are experienced with transferring and applying knowledge from the academic field to the system and application communities
- Interested groups working on EU projects from various fields related to MLIA/CLIR
- Selected representatives of digital libraries and information providers that are concerned with multilingual access to their data.

The number of attendees of the workshop was restricted, in order to get a lively and focused discussion. In fact, all attendees provided input, mostly through presentations and subsequent

discussions, which will be used as source for future analysis by the TrebleCLEF consortium. Although the workshop included participants with varying academic and commercial interests, all attendees shared a background of (sometimes limited) involvement with CLEF (either as organisers or participants) and could therefore provide helpful insights both from the outside and inside on the use of CLEF results within an operational context.

### 3 Goals

The outputs of this workshop form one of the key sources in preparing best practice recommendations that will be published as project Deliverable 3.3. In order to support the creation of this report, the following goals for the workshop were communicated to the attendees in advance:

- To discuss needs when adapting systems, such as digital libraries and information portals, for multilingual access
- To discuss the challenges in transferring the assembled academic work on MLIA/CLIR for use by system developers/implementers and to detect open issues
- To identify best practices in MLIA system development/implementation
- To present to the attendees the output of a similar workshop focussed on the user requirements of MLIA/CLIR held in Segovia/Spain in June 08 [4].

The workshop programme was initially conceived to begin with a series of presentations by the attendees, which would detail their experiences with MLIA systems in operational settings. These presentations would form the basis for a series of break-out sessions, where small groups could formulate questions and/or challenges related to MLIA research and the transfer of MLIA technology to operational settings. The output of these break-out sessions would then form the basis of a plenary discussion, highlighting the pressing issues in MLIA system development, and agreeing on recommendations for MLIA system developers.

However, it became apparent early on that discussion during and after the presentations was much deeper and more intensive than anticipated. Input was also of very high quality, and therefore required more time to discuss the issues in adequate detail. A far more substantial set of questions and challenges came directly from these discussions than initially expected and it was therefore agreed (with attendees) to expand time allocated for both the presentations and associated discussion. This meant that the planned break-out sessions did not take place as expected. The output produced during these discussions proved to be very suitable for the subsequent plenary session in the workshop, and we are confident that the results from the workshop greatly benefitted from this decision.

### 4 Presentations

The following presentations were given during the workshop:

- "Welcome" by Hans-Peter Hutter and Martin Braschler, Zurich Univ. of Applied Sciences
  - The presentation gave a short overview of the host organization ZHAW and the organization of the workshop
- "What happened at CLEF 2008?" by Paul Clough, University of Sheffield [5]
  - Paul Clough presented an overview of the CLEF 2008 campaign that had culminated with the CLEF 2008 workshop in Aarhus, Denmark a few weeks earlier. His overview

was mainly designed to ensure that all attendees have the full details of current work in the CLEF campaign context.

- "Successful blueprint for CLEF participation - but does it have any relation to reality?" by Martin Braschler [6]
  - Martin Braschler presented preliminary work on analysis of experiments by different CLEF participants mainly in the ad-hoc track. This analysis seems to indicate that there is a set of well-known "tweaks" that allow one to obtain high scores as defined in the CLEF evaluation methodology. The presentation served as a starting point for discussion on the validity of these tweaks in operational settings
- "What happened at the Segovia User Workshop?" by Martin Braschler with material provided by Julio Gonzalo, UNED [7]
  - A short overview of the output of the Segovia TrebleCLEF user workshop was presented to the attendees. The Segovia workshop together with the Winterthur system developers workshop both serve as inputs into the forthcoming TrebleCLEF deliverable 3.3.
- "User needs studies and surveys: MLIA needs in the cultural heritage domain" by Paul Clough [8]
  - Paul Clough presented work carried out at University of Sheffield on studying user's attitude towards MLIA in the cultural heritage domain. They were able to carry out an online survey among users of the Tate Gallery's web site. This was supplemented with further insights collected by interviewing cultural heritage professionals. He pointed at issues still remaining with acquiring suitable resources, and finding a form of presentation of foreign-language results that is helpful to the users.
- "Multilingual Information Access – in the lab and in the wild" by Jochen Leidner, ThomsonReuters [9]
  - Jochen Leidner presented a critical view of MLIA and the state of the academic field, presenting his thoughts on the pressing issues and possible approaches. He points out that there are still uncertainties about use cases and business cases associated with MLIA. Yet, he is convinced that MLIA and cross-language search are very important, as evidenced with progressing globalization of business and politics. Main unsolved issues for him include: systems remain complex and expensive (especially when many languages are involved), machine translation is still a problematic solution for document presentation, evaluation of MLIA systems presently controls mostly the test collection, and less the system parameters, and important language resources are not widely available. He is convinced that CLEF and TrebleCLEF are and can continue to be very helpful to address these issues.
- "CLIR information retrieval in digital libraries" by Luca Dini, CELI [10]
  - Luca Dini presented insights from his work on MLIA in the context of digital libraries. He points out that three quarters of users want to search for books by free-text. Multilingual mapping of subject headings is not an ideal solution in this case. He points to vocabulary coverage by the translation resources as one of the most serious issues encountered in his work. In free-text search, infrequent search terms are often crucial. A number of approaches to this problem were presented. Other major problems mentioned by Luca Dini were semantic ambiguity and the handling of named entities. Finally, the CACAO service and its evaluation at CLEF were presented.
- "TrebleCLEF" by Stephane Clinchant, XRCE Xerox Grenoble [11]
  - Stephane Clinchant contrasted MLIA with monolingual information retrieval. He pointed out that the notion of "relevance" probably needs to be expanded with one of

"understandability" when assessing the results of an MLIA system. This understandability, and thus the best resources/techniques for producing the multilingual search result, may depend on the user. As Luca Dini before him, Stephane Clinchant also highlighted the problem of vocabulary coverage in translation resources, pointing to a precision/coverage conflict when compiling such resources. He then proceeded to present some examples of work at XRCE, pointing out the difficulties that were encountered.

- "2 examples of real applications of cross-language information access" by Christian Fluhr, NewPhenix [12]
  - As stated in the title of his presentation, Christian Fluhr discussed insights from his involvement in two specific projects: cross-language access to an OPAC, and cross-language access for companies selling photos over the internet. He showed how morphological analysis was successfully paired with bilingual dictionaries to produce better transfers between source and target languages. Combined with a handling of term dependencies, this allows the generation of domain-specific translations. He made a good case for multilingual access in libraries, where accessible content representations are typically very short. In such cases, all available data needs to be used, regardless of language. Similar problems are found when dealing with captions describing photographs; in addition, some photographers have a tendency to "stuff" the list of keywords associated with a picture in order to maximize sales. Search results are improved by using a feedback technique, with the end user selecting good examples from an initial result list.
- "DRIVER Repository Infrastructure" by Paolo Manghi, ISTI-CNR [13]
  - Paolo Manghi gave a short presentation of the DRIVER EU Project. The goal is to provide monolingual and multilingual functionality built as a set of services. These services are configured with respect to the languages that they process ("language kits"). Presently, the DRIVER consortium is harvesting open access research publications, and has collected over 700,000 publications in more than 40 languages.
- "MultiMatch: Multilingual/Multimedia Access to Cultural Heritage" by Franca Debole, ISTI-CNR [14]
  - Franca Debole gave a short overview of the MultiMatch EU FP6 project. The project deals with multilingual and multimedia access to cultural heritage content. She gave a summary of the challenges that were encountered when handling a range of languages, such as widely varying availability and quality of resources. In addition, she provided insight in the adaptation of a commercial machine translation system to their MLIA system.
- "MLIA and Digital Library Systems" by Nicola Ferro, Univ. Padova [15]
  - Nicola Ferro provided an overview of MLIA-related work conducted in the context of the TEL ("The European Library") system and the "Europeana" project. He presented the Europeana Object Model and discussed the implications of using surrogates to represent information objects in the MLIA context. CLEF 2008 has hosted a "TEL" track in collaboration with The European Library, and Ferro gave an overview of the data used in this track, and the experiments that were conducted on this data. In the future, it is envisioned to conduct more Grid experiments in the CLEF campaigns, and a short overview of how this will be setup was provided. The main message by Ferro was that MLIA should not be "tacked on", but instead be designed as an integral part of the system architecture and object model.

## 5 Discussions

### 5.1 Questions

As the immediate result of the discussions accompanying the presentations, a list of "open questions" was compiled during the workshop itself. This list of questions was then debated at the plenary session on Friday afternoon. For easier discussion, the questions were subdivided into four categories: (1) CLEF tasks, (2) Resources for MLIA, (3) Components for MLIA, and (4) Varia. Cleaned up lists of the questions that were debated at the plenary session are provided below.

#### Category 1: CLEF tasks

This category contains the questions raised with respect to the organization of the CLEF tasks, and the experiments that are carried out by CLEF participants during the individual evaluation campaigns. Questions that were identified for this category and discussed at the plenary were:

- CLEF is structured into tracks and tasks. Which of these tasks are realistic, and which ones are artificial?
- Is there a recipe on "How to win CLEF", i.e. how to maximize effectiveness measures as used by CLEF? Is this recipe generalizable and relevant beyond the CLEF campaigns?
- Which widely adopted tools in CLEF experiments (stemming, stopwording, relevance feedback etc.) are useful in operational settings? Which of these tools are acceptable to real users?
- What part of the data should be translated in operational settings? Queries or documents? Can we learn this from CLEF experiments?
- What should be measured in CLEF experiments? System effectiveness? Other aspects?
- How should CLEF queries be constructed? How can they be made relevant for testing of operational systems
- Are the CLEF queries too artificial? (too clean, not interpretable)
- How can we account for user interactivity or query refinement? (CLEF experiments tend to be "one-off", query/result experiments, not investigating further interaction of the user based on the initial result)
- Is there a "template" on how to define a new evaluation task?

#### Category 2: Resources for MLIA

This category contains the questions raised with respect to the availability/suitability/creation etc. of resources for MLIA system development. Questions that were identified for this category and discussed at the plenary were:

- What is the potential for standardization of MLIA resources? Can we generalize the build process or other aspects?
- What is the "critical mass" for MLIA resources? I.e., how many translation pairs do we need for a dictionary? How large does the vocabulary coverage have to be?
- How do we combine resources (of the same type or of different types)?

- How can we acquire resources for less widely spoken languages, where limited digital content is readily available?
- What is the trade-off between precision of entries in the MLIA resources versus a maximization of vocabulary coverage?
- When is the use of pivot languages acceptable? That is, can we use an intermediary language where direct translation is not feasible?
- Can we "extract" resources from openly available data sources like Wikipedia?
- How do we deal with community-constructed resources? This includes resources constructed by interaction between the end users and the system (i.e. by allowing the end users to suggest alternative translations)

### **Category 3: Components for MLIA**

This category contains the questions raised with respect to the building and deployment of MLIA system components. Questions that were identified for this category and discussed at the plenary were:

- Is there a suitable framework for MLIA system components? Can we generalize this to multiple settings?
- Are there "standard flows" for MLIA that can be implemented by "chaining" these components?
- How can components be used to upgrade monolingual information access systems?
- How can MLIA functionality be packaged into services? How should the interfaces for these services be designed?
- Is it possible to package "language kits", containing the basic components to enhance a system for a specific language? How can these kits be bootstrapped for new languages?
- How do we encourage the sharing of components? How do we enable easier adaptability of such components?
- What are the implications of frequent updates to components (or resources)?
- How much control over components is necessary?

### **Category 4: Varia**

This category summarizes all other miscellaneous questions that were raised during the presentations:

- How can we sell MLIA as a "feature" to service providers? What are the benefits? What are the costs?
- Are there feasible "low-budget" solutions?
- It is often claimed that "Machine translation does not work" – how should results of MLIA be presented?
- What are real-world use cases for MLIA? What is the use of presenting results that the searcher cannot read/interpret?
- Are there real-world document collections for MLIA available for testing?
- How can we replicate experiments carried out by CLEF/MLIA researchers?
- Should search results be more question-driven or more user-driven?

- How do users cope with ambiguity introduced by the translation step? How should we handle this problem in the systems?
- Can online services be used as replacements for certain resources in operations systems? What is the reliability if there is no direct control over these resources?
- What if documents are misleading? Examples include spam on the internet or misleading information.
- How much "natural language" interpretation of queries, such as morphological analysis, part-of-speech tagging and word sense disambiguation, is necessary?
- How do we properly present the results? Should this be different from monolingual settings? Are clustering, snippets, multilingual result lists helpful?
- How should we aggregate multilingual information?
- What is the relationship of multilanguage and multimedia?
- Is there a need for more/better open source components? Such as more intensive development of open source machine translation systems?

## 5.2 Category CLEF tasks

It was argued in the plenary session that one of the main challenges with respect to transfer of CLEF results to an operational setting is the successful replication of experiments that are reported in papers written by CLEF participant. The attendees expressed regret that many papers do not present results in a "hypothesis-driven" way and is often unclear what research hypotheses are being tested. Instead papers focus on tabulating different evaluation figures. This incurs the danger that practitioners in search of a solution or blueprint for MLIA issues find it hard to match their needs to the relevant experiment descriptions.

Should this lack of focus on hypothesis testing prove true – and it has to be stressed that while this feeling was shared by many workshop participants no exhaustive analysis of the CLEF paper corpus in this regard has been undertaken so far – this also poses a danger with respect to the replicability of the scientific experiments - one of the cornerstones of good, sustainable evaluation. Indeed, anecdotal evidence for problems in this regard and with respect to precision in reporting system descriptions was mentioned by workshop participants.

The question of whether "successful" CLEF approaches could be generalised was also discussed. The workshop attendees considered many of the achievements of CLEF participants to be highly relevant to technology transfer. As evidence for this the growing amount of performance data for different weighting schemes assembled over the different years of CLEF testing was cited. In this sense, the compilation of best-practice recommendations based on successful CLEF experiments was very welcomed by the workshop participants – and is indeed the goal of Work Package 3.

The hope that an aggregation of system descriptions may help to uncover gaps in the coverage of the CLEF campaigns was also expressed. Some limited work in this respect will also be undertaken for Deliverable 3.3. Areas which were identified as interesting for future CLEF campaign experimentation with respect to MLIA in operational systems included:

- More research on result list fusion (termed a "bottleneck" by workshop attendees)
- How to filter results lists with respect to "back translation" of queries and documents for presentation to the user (i.e., an "understandability" criterion)

- More performance measurements
- Larger sets of queries (>100 queries per task; potentially assembled over multiple campaigns)
- "Coded queries" instead of full-text queries
- Grid experiments (note: this is presently addressed by TrebleCLEF)
- Identification of solid blueprints for MLIA

### 5.3 Resources for MLIA

Discussion of issues with respect to "Resources for MLIA" centred on multiple issues. The workshop participants were mainly interested in issues that are, broadly speaking, related to resource maintenance. How to evaluate and handle uneven quality of resources? How to adapt resources for individual systems/setup? What is the "critical mass" for vocabulary coverage? It was felt that the adaptability of (translation) resources is key – in many settings even users must be able to modify resources. Furthermore, it was pointed out that it is highly beneficial to (hand-) tune resources to cover the most frequent information need of individual systems.

### 5.4 Components for MLIA

The workshop attendees showed great interest in questions surrounding "components for MLIA". It was felt that there is a large untapped potential in working on the breaking down of MLIA systems and processes into components and then trying to build frameworks or describe processes that use them. There is currently little work on "MLIA software engineering" in CLEF. Workshop attendees showed great interest in work on an MLIA reference model.

Successfully defining components will make their individual evaluation more important. Can some components be isolated for this purpose, or are they linked too closely (i.e. influence each other substantially)?

### 5.5 Varia

Discussion of the questions formulated in the "varia" category centred on results presentation, and involvement of the user in the translation process. It was unclear to workshop attendees as to what degree system users can be involved in the technical necessities of query translation. Specifically, can users understand/interpret "back-translations" of queries, e.g. the English translation of a German rendering of an English original query? This issue is relevant in case that the initial German translation of the English query is ambiguous; in such a case, interpretation of search results can be very difficult for users if the same ambiguity does not exist in the source language.

Workshop attendees also expressed interest in more research/work on result presentation, specifically the use of "snippets" to represent documents in the result lists. In the case of MLIA, these snippets potentially represent documents that are written in a language different from the source language. The problem is the generation of good document summaries in the source query language, i.e. translated from the document content.

## 6 Summary

The workshop was designed to provide input for the preparation of Deliverable 3.3, "System-oriented and user-oriented MLIA best-practice recommendations". The workshop has been very successful in

compiling lists of open issues in MLIA system development, and in discussing approaches to address some of these issues. In this sense, we hope to use the input of the workshop's attendees to facilitate the compilation of best practice recommendations in the deliverable.

We were delighted to hear from the workshop's attendees that the goals of Deliverable 3.3 are indeed very much anticipated: the attendees identified the lack of best practice recommendations in MLIA as a problem, and indicated that an abstraction of successful blueprints compiled from the experiences by CLEF participants, such as planned for Deliverable 3.3, would be a very welcome help for practitioners. Attendees also expressed interest in the aggregation of numbers for different usages of individual components and techniques by CLEF participants, giving an indication on how widely such components and techniques are adopted and tested. Again, it is planned to address this issue in deliverable 3.3 to such a degree as is possible by semi-automatic analysis of the CLEF participant's system descriptions.

## Acknowledgements

We thank the attendees of the TrebleCLEF Systems Developers Workshop for their presentations and participation in the discussions.

## References

- [1] TrebleCLEF homepage, <http://www.trebleclef.eu>
- [2] CLEF homepage, <http://www.clef-campaign.org>
- [3] TrebleCLEF System Developers Workshop homepage, <http://www.init.zhaw.ch/trebleclef/workshop>
- [4] TrebleCLEF User Workshop: Workshop on Best Practices for the development of Multilingual Information Access systems: The User Perspective, <http://nlp.uned.es/trebleclef-user-workshop/>
- [5] Clough, P.: "What happened at CLEF 2008?" Presentation at the TrebleCLEF Systems Developers Workshop. See [3]
- [6] Braschler, M.: "A Successful Blueprint for CLEF Participation". Presentation at the TrebleCLEF Systems Developers Workshop. See [3]
- [7] Gonzalo, J. and Braschler, M.: "What happened at the Segovia User Workshop?" Presentation at the TrebleCLEF Systems Developers Workshop. See [3]
- [8] Clough, P.: "User Needs Studies and Surveys: MLIA Needs in the Cultural Heritage Domain" Presentation at the TrebleCLEF Systems Developers Workshop. See [3]
- [9] Leidner, J.: "Multilingual Information Access – in the Lab and in the Wild" Presentation at the TrebleCLEF Systems Developers Workshop. See [3]
- [10] Dini, L.: "CLIR Information Retrieval in Digital Libraries" Presentation at the TrebleCLEF Systems Developers Workshop. See [3]
- [11] Renders, J.M. and Clinchant, S.: "TrebleCLEF" Presentation at the TrebleCLEF Systems Developers Workshop. See [3]
- [12] Fluhr, C.: "2 Examples of Real Applications of Cross-Language Information Access" Presentation at the TrebleCLEF Systems Developers Workshop. See [3]
- [13] Manghi, P.: "DRIVER Repository Infrastructure" Presentation at the TrebleCLEF Systems Developers Workshop. See [3]
- [14] Debole, F.: "MultiMatch: Multilingual/Multimedia Access to Cultural Heritage" Presentation at the TrebleCLEF Systems Developers Workshop. See [3]
- [15] Ferro, N.: "MLIA and Digital Library Systems" Presentation at the TrebleCLEF Systems Developers Workshop. See [3]

## **Appendix – Attendees (alphabetic listing)**

- Martin Braschler, InIT, Zurich University of Applied Sciences
- Stephane Clinchant, Xerox XRCE Grenoble
- Paul Clough, University of Sheffield
- Franca Debole, ISTI-CNR
- Luca Dini, CELI
- Nicola Ferro, University of Padova
- Christian Fluhr, New Phenix
- Hans-Petter Hutter, InIT, Zurich University of Applied Sciences
- Jochen Leidner, Thomson Reuters
- Loic Maisonnasse, IMAG
- Paolo Manghi, ISTI-CHR
- Natalia Manola, Dept. Informatics, University of Athens
- Jacques Savoy, Universite de Neuchatel