

# CLEF 2007: Ad Hoc Track Overview

Giorgio M. Di Nunzio<sup>1</sup>, Nicola Ferro<sup>1</sup>, Thomas Mandl<sup>2</sup>, and Carol Peters<sup>3</sup>

<sup>1</sup> Department of Information Engineering, University of Padua, Italy

{dinunzio,ferro}@dei.unipd.it

<sup>2</sup> Information Science, University of Hildesheim, Germany

mandl@uni-hildesheim.de

<sup>3</sup> ISTI-CNR, Area di Ricerca, Pisa, Italy

carol.peters@isti.cnr.it

**Abstract.** We describe the objectives and organization of the CLEF 2007 Ad Hoc track and discuss the main characteristics of the tasks offered to test monolingual and cross-language textual document retrieval systems. The track was divided into two streams. The main stream offered mono- and bilingual tasks on target collections for central European languages (Bulgarian, Czech and Hungarian). Similarly to last year, a bilingual task that encouraged system testing with non-European languages against English documents was also offered; this year, particular attention was given to Indian languages. The second stream, designed for more experienced participants, offered mono- and bilingual “robust” tasks with the objective of privileging experiments which achieve good stable performance over all queries rather than high average performance. These experiments re-used CLEF test collections from previous years in three languages (English, French, and Portuguese). The performance achieved for each task is presented and discussed.

## 1 Introduction

The Ad Hoc retrieval track is generally considered to be the core track in the *Cross-Language Evaluation Forum (CLEF)*. The aim of this track is to promote the development of monolingual and cross-language textual document retrieval systems. Similarly to last year, the CLEF 2007 ad hoc track was structured in two streams. The main stream offered mono- and bilingual retrieval tasks on target collections for central European languages plus a bilingual task encouraging system testing with non-European languages against English documents. The second stream, designed for more experienced participants, was the “robust task”, aimed at finding documents for very difficult queries. It used test collections developed in previous years.

The **Monolingual** and **Bilingual** tasks were principally offered for Bulgarian, Czech and Hungarian target collections. Additionally, a bilingual task was offered to test querying with non-European language queries against an English target collection. As a result of requests from a number of Indian research institutes, a special sub-task for Indian languages was offered with topics in Bengali, Hindi,

Marathi, Tamil and Telugu. The aim in all cases was to retrieve relevant documents from the chosen target collection and submit the results in a ranked list.

The **Robust** task proposed mono- and bilingual experiments using the test collections built over the last six CLEF campaigns. Collections and topics in English, Portuguese and French were used. The goal of the robust analysis is to improve the user experience with a retrieval system. Poor performing topics are more serious for the user than performance losses in the middle and upper interval. The robust task gives preference to systems which achieve a minimal level for all topics. The measure used to ensure this is the geometric mean over all topics. The robust task intends to evaluate stable performance over all topics instead of high average performance. Experiments are offered with a larger topic set.

This was the first year since CLEF began that we have not offered a **Multilingual** ad hoc task (ie searching a target collection in multiple languages).

In this paper we describe the track setup, the evaluation methodology and the participation in the different tasks (Section 2), present the main characteristics of the experiments and show the results (Sections 3 - 5). The final section provides a brief summing up. For information on the various approaches and resources used by the groups participating in this track and the issues they focused on, we refer the reader to the other papers in the Ad Hoc section of these Proceedings.

## 2 Track Setup

The Ad Hoc track in CLEF adopts a corpus-based, automatic scoring method for the assessment of system performance, based on ideas first introduced in the Cranfield experiments in the late 1960s. The test collection used consists of a set of “topics” describing information needs and a collection of documents to be searched to find those documents that satisfy these information needs. Evaluation of system performance is then done by judging the documents retrieved in response to a topic with respect to their relevance, and computing the recall and precision measures. The distinguishing feature of CLEF is that it applies this evaluation paradigm in a multilingual setting. This means that the criteria normally adopted to create a test collection, consisting of suitable documents, sample queries and relevance assessments, have been adapted to satisfy the particular requirements of the multilingual context. All language dependent tasks such as topic creation and relevance judgment are performed in a distributed setting by native speakers. Rules are established and a tight central coordination is maintained in order to ensure consistency and coherency of topic and relevance judgment sets over the different collections, languages and tracks.

### 2.1 Test Collections

Different test collections were used in the ad hoc task this year. The main stream used national newspaper documents from 2002 as the target collections, creating sets of new topics and making new relevance assessments. The robust task reused existing CLEF test collections and did not create any new topics or make any fresh relevance assessments.

**Table 1.** Test collections for the main stream Ad Hoc tasks

Language	Collections
Bulgarian	Sega 2002, Standart 2002, Novinar 2002
Czech	Mlada fronta DNES 2002, Lidové Noviny 2002
English	LA Times 2002
Hungarian	Magyar Hirlap 2002

**Table 2.** Test collections for the Robust task

Language	Collections
English	LA Times 94, Glasgow Herald 95
French	ATS (SDA) 94/95, Le Monde 94
Portuguese	Publico 94/95, Folha de Sao Paulo 94/95

**Documents.** The document collections used for the CLEF 2007 Ad Hoc tasks are part of the CLEF multilingual corpus of newspaper and news agency documents described in the Introduction to these Proceedings.

In the main stream monolingual and bilingual tasks, Bulgarian, Czech, Hungarian and English national newspapers for 2002 were used. Much of this data represented new additions to the CLEF multilingual comparable text corpora: Czech is a totally new language in the ad hoc track although it was introduced into the speech retrieval track last year; the Bulgarian collection was expanded with the addition of another national newspaper, and in order to have comparable data for English, we acquired a new American-English collection: Los Angeles Times 2002. Table 1 summarizes the collections used for each language.

The robust task used test collections containing news documents for the period 1994-1995 in three languages (English, French, and Portuguese) used in CLEF 2000 through CLEF 2006. Table 2 summarizes the collections used for each language.

**Topics.** Topics in the CLEF ad hoc track are structured statements representing information needs; the systems use the topics to derive their queries. Each topic consists of three parts: a brief “title” statement; a one-sentence “description”; a more complex “narrative” specifying the relevance assessment criteria.

Sets of 50 topics were created for the CLEF 2007 ad hoc mono- and bilingual tasks. All topic sets were created by native speakers. One of the decisions taken early on in the organization of the CLEF ad hoc tracks was that for each task the same set of topics, rendered in different languages, would be used to query the different collections. There were a number of reasons for this: it makes it easier to compare results over different target languages, it means that there is a single master set that is rendered in all query languages, and a single set of relevance assessments for each language is sufficient for all tasks. In CLEF 2006 we deviated from this rule as we were using document collections from two distinct periods (1994/5 and 2002) and created partially separate (but overlapping) sets with a common set of time-independent topics and separate sets of time-specific topics. As we had expected this really complicated our lives as we had to build more

topics and had to specify very carefully which topic sets were to be used against which document collections<sup>1</sup>. We decided not to repeat this experience this year and thus only used collections from the same time period.

We created topics in both European and non-European languages. European language topics were offered for Bulgarian, Czech, English, French, Hungarian, Italian and Spanish. The non-European topics were prepared according to demand from participants. This year we had Amharic, Chinese, Indonesian, Oromo plus the group of Indian languages: Bengali, Hindi, Marathi, Tamil and Telugu.

The provision of topics in unfamiliar scripts did lead to some problems. These were not caused by encoding issues (all CLEF data is encoded using UTF-8) but rather by errors in the topic sets which were very difficult for us to spot. Although most such problems were quickly noted and corrected, and the participants were informed so that they all used the right set, one did escape our notice: the title of Topic 430 in the Czech set was corrupted and systems using Czech thus did not do well with this topic. It should be remembered, however, that an error in one topic does not really impact significantly on the comparative results of the systems. The topic will, however, be corrected for future use.

This year topics have been identified by means of a Digital Object Identifier (DOI)<sup>2</sup> of the experiment [1] which allows us to reference and cite them. Below we give an example of the English version of a typical CLEF 2007 topic:

```
<top lang="en">
<num>10.2452/401-AH</num>
<title>Euro Inflation</title>
<desc>Find documents about rises in prices after the introduction of the
Euro.</desc>
<narr>Any document is relevant that provides information on the rise of
prices in any country that introduced the common European
currency.</narr>
</top>
```

For the robust task, topic sets from CLEF 2001 to 2006 in English, French and Portuguese were used. For English and French, in CLEF for more time, training topics were offered and a set of 100 topics were used for testing. For Portuguese, no training topics were possible and a set of 150 test topics was used.

## 2.2 Participation Guidelines

To carry out the retrieval tasks of the CLEF campaign, systems have to build supporting data structures. Allowable data structures include any new structures built automatically (such as inverted files, thesauri, conceptual networks, etc.) or manually (such as thesauri, synonym lists, knowledge bases, rules, etc.) from the documents. They may not, however, be modified in response to the topics,

<sup>1</sup> This is something that anyone reusing the CLEF 2006 ad hoc test collection needs to be very careful about.

<sup>2</sup> In order to resolve the DOIs used in this paper and to access on-line the relative information, you can use any DOI resolver, such as <http://dx.doi.org/>.

e.g. by adding topic words that are not already in the dictionaries used by their systems in order to extend coverage.

Some CLEF data collections contain manually assigned, controlled or uncontrolled index terms. The use of such terms is limited to specific experiments that have to be declared as “manual” runs.

Topics can be converted into queries that a system can execute in many different ways. CLEF strongly encourages groups to determine what constitutes a base run for their experiments and to include these runs (officially or unofficially) to allow useful interpretations of the results. Unofficial runs are those not submitted to CLEF but evaluated using the `trec_eval` package. This year we have used the new package written by Chris Buckley for the *Text REtrieval Conference (TREC)* (`trec_eval` 8.0) and available from the TREC website<sup>3</sup>.

As a consequence of limited evaluation resources, a maximum of 12 runs each for the mono- and bilingual tasks was allowed (no more than 4 runs for any one language combination - we try to encourage diversity). For bi- and monolingual robust tasks, 4 runs were allowed per language or language pair.

### 2.3 Relevance Assessment

The number of documents in large test collections such as CLEF makes it impractical to judge every document for relevance. Instead approximate recall values are calculated using pooling techniques. The results submitted by the groups participating in the ad hoc tasks are used to form a pool of documents for each topic and language by collecting the highly ranked documents from selected runs according to a set of predefined criteria. Traditionally, the top 100 ranked documents from each of the runs selected are included in the pool; in such a case we say that the pool is of depth 100. This pool is then used for subsequent relevance judgments. After calculating the effectiveness measures, the results are analyzed and run statistics produced and distributed.

The stability of pools constructed in this way and their reliability for post-campaign experiments is discussed in [2] with respect to the CLEF 2003 pools. New pools were formed in CLEF 2007 for the runs submitted for the main stream mono- and bilingual tasks. Instead, the robust tasks used the original pools and relevance assessments from previous CLEF campaigns.

The main criteria used when constructing these pools were:

- favour diversity among approaches adopted by participants, according to the descriptions of the experiments provided by the participants;
- choose at least one experiment for each participant in each task, selected from the experiments with highest priority as indicated by the participant;
- add mandatory title+description experiments, even though they do not have high priority;
- add manual experiments, when provided;
- for bilingual tasks, ensure that each source topic language is represented.

---

<sup>3</sup> [http://trec.nist.gov/trec\\_eval/](http://trec.nist.gov/trec_eval/)

One important limitation when forming the pools is the number of documents to be assessed. We estimate that assessors can judge from 60 to 100 documents per hour, providing binary judgments: relevant / not relevant. This is actually an optimistic estimate and shows what a time-consuming and resource expensive task human relevance assessment is. This limitation impacts strongly on the application of the criteria above - and implies that we are obliged to be flexible in the number of documents judged per selected run for individual pools.

This meant that this year, in order to create pools of more-or-less equivalent size (approx. 20,000 documents), the depth of the Bulgarian, Czech and Hungarian pools varied: 60 for Czech and 80 for Bulgarian and Hungarian, rather than the depth of 100 originally used to judge TREC ad hoc experiments<sup>4</sup>. In his paper in these proceedings, Tomlinson [3] makes some interesting observations in this respect. He claims that on average, the percentage of relevant items assessed was less than 60% for Czech, 70% for Bulgarian and 85% for Hungarian. However, as Tomlinson also points out, it has already been shown that test collections created in this way do normally provide reliable results, even if not all relevant documents are included in the pool.

When building the pool for English, in order to respect the above criteria and also to obtain a pool depth of 60, we had to include more than 25,000 documents. Even so, as can be seen from Table 3, it was impossible to include very many runs - just one monolingual and one bilingual run for each set of experiments.

The box plot of Figure 1 compares the distributions of the relevant documents across the topics of each pool for the different ad hoc pools; the boxes are ordered by decreasing mean number of relevant documents per topic. As can be noted, Bulgarian, Czech, and Hungarian distributions appear similar, even though the Czech and Hungarian ones are slightly more asymmetric towards topics with a greater number of relevant documents. On the other hand, the English distribution presents a greater number of relevant documents per topic, with respect to the other distributions, and is quite asymmetric towards topics with a greater number of relevant documents. All the distributions show some upper outliers, i.e. topics with a great number of relevant document with respect to the behaviour of the other topics in the distribution. These outliers are probably due to the fact that CLEF topics have to be able to retrieve relevant documents in all the collections; therefore, they may find considerably more relevant documents in one collection than in others depending on the contents of the separate datasets. Thus, typically, each pool will have a different set of outliers.

Table 3 reports summary information on the 2007 ad hoc pools used to calculate the results for the main monolingual and bilingual experiments. In particular, for each pool, we show the number of topics, the number of runs submitted, the number of runs included in the pool, the number of documents in the pool (relevant and non-relevant), and the number of assessors.

---

<sup>4</sup> Tests made on NTCIR pools in previous years have suggested that a depth of 60 is normally adequate to create stable pools, presuming that a sufficient number of runs from different systems have been included.

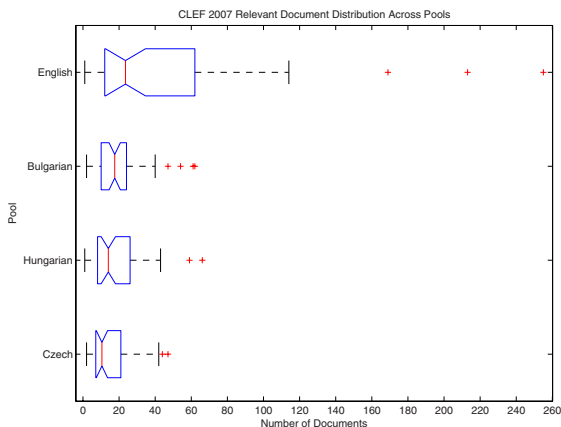


Fig. 1. Distribution of the relevant documents across the pools

## 2.4 Result Calculation

Evaluation campaigns such as TREC and CLEF are based on the belief that the effectiveness of *Information Retrieval Systems (IRSs)* can be objectively evaluated by an analysis of a representative set of sample search results. For this, effectiveness measures are calculated based on the results submitted by the participants and the relevance assessments. Popular measures usually adopted for exercises of this type are Recall and Precision. Details on how they are calculated for CLEF are given in [5]. For the robust task, we used different measures, see below Section 5.

The individual results for all official ad hoc experiments in CLEF 2007 are given in the Appendixes at the end of the Working Notes distributed for the workshop [6,7].

## 2.5 Participants and Experiments

22 groups from 12 different countries submitted results for one or more of the ad hoc tasks - a slight decrease on the 25 participants of last year. These groups submitted a total of 235 runs, a decrease of about 20% on the 296 runs of 2006. The average number of submitted runs per participant also slightly decreased: from 11.7 runs/participant of 2006 to 10.6 runs/participant of this year.

Participants were required to submit at least one title+description (“TD”) run per task in order to increase comparability between experiments. The large majority of runs (138 out of 235, 58.72%) used this combination of topic fields, 50 (21.28%) used all fields, 46 (19.57%) used the title field, and only 1 (0.43%) used just the description field. The majority of experiments were conducted using automatic query construction (230 out of 235, 97.87%) and only in a small fraction of the experiments (5 out 237, 2.13%) were queries been manually constructed from topics. A breakdown into the separate tasks is shown in Table 4(a).

**Table 3.** Summary information about CLEF 2007 pools

<b>Bulgarian Pool</b>	
<b>Pool size</b>	19,441 pooled documents <ul style="list-style-type: none"> <li>– 18,429 not relevant documents</li> <li>– 1,012 relevant documents</li> </ul> 50 topics
<b>Pooled Experiments</b>	13 out of 18 submitted experiments <ul style="list-style-type: none"> <li>– monolingual: 11 out of 16 submitted experiments</li> <li>– bilingual: 2 out of 2 submitted experiments</li> </ul>
<b>Assessors</b>	4 assessors
<b>Czech Pool</b>	
<b>Pool size</b>	20,607 pooled documents <ul style="list-style-type: none"> <li>– 19,485 not relevant documents</li> <li>– 762 relevant documents</li> </ul> 50 topics
<b>Pooled Experiments</b>	19 out of 29 submitted experiments <ul style="list-style-type: none"> <li>– monolingual: 17 out of 27 submitted experiments</li> <li>– bilingual: 2 out of 2 submitted experiments</li> </ul>
<b>Assessors</b>	4 assessors
<b>English Pool</b>	
<b>Pool size</b>	24,855 pooled documents <ul style="list-style-type: none"> <li>– 22,608 not relevant documents</li> <li>– 2,247 relevant documents</li> </ul> 50 topics
<b>Pooled Experiments</b>	20 out of 104 submitted experiments <ul style="list-style-type: none"> <li>– monolingual: 10 out of 31 submitted experiments</li> <li>– bilingual: 10 out of 73 submitted experiments</li> </ul>
<b>Assessors</b>	5 assessors
<b>Hungarian Pool</b>	
<b>Pool size</b>	18,704 pooled documents <ul style="list-style-type: none"> <li>– 17,793 not relevant documents</li> <li>– 911 relevant documents</li> </ul> 50 topics
<b>Pooled Experiments</b>	14 out of 21 submitted experiments <ul style="list-style-type: none"> <li>– monolingual: 12 out of 19 submitted experiments</li> <li>– bilingual: 2 out of 2 submitted experiments</li> </ul>
<b>Assessors</b>	6 assessors



**Table 4.** Breakdown of experiments into tracks and topic languages

(a) Number of experiments per track, participant.

Track	# Part.	# Runs
Monolingual-BG	5	16
Monolingual-CS	8	27
Monolingual-EN	10	31
Monolingual-HU	6	19
Bilingual-X2BG	1	2
Bilingual-X2CS	1	2
Bilingual-X2EN	10	73
Bilingual-X2HU	1	2
Robust-Mono-EN	3	11
Robust-Mono-FR	5	12
Robust-Mono-PT	4	11
Robust-Bili-X2FR	3	9
Robust-Training-Mono-EN	2	6
Robust-Training-Mono-FR	2	6
Robust-Training-Bili-X2FR	2	8
<b>Total</b>		<b>235</b>

(b) List of experiments by topic language.

Topic Lang.	# Runs
English	73
Hungarian	33
Czech	26
Bulgarian	16
Indonesian	16
French	14
Hindi	13
Chinese	12
Portuguese	11
Amharic	9
Bengali	4
Oromo	4
Marathi	2
Telugu	2
<b>Total</b>	<b>235</b>

Fourteen different topic languages were used in the ad hoc experiments. As always, the most popular language for queries was English, with Hungarian second. The number of runs per topic language is shown in Table 4(b).

### 3 Main Stream Monolingual Experiments

Monolingual retrieval focused on central-European languages this year, with tasks offered for Bulgarian, Czech and Hungarian. Eight groups presented results for 1 or more of these languages. We also requested participants in the bilingual-to-English task to submit one English monolingual run, but only in order to provide a baseline for their bilingual experiments and in order to strengthen the English pool for relevance assessment<sup>5</sup>.

Five of the participating groups submitted runs for all three languages. One group was unable to complete its Bulgarian experiments, submitting results for just the other two languages. However, they subsequently completed their work on Bulgarian post-campaign, and results for all three languages are reported in this volume [11]. The two groups from the Czech Republic only submitted runs for Czech. From Table 5, it can be seen that the best performing groups were more-or-less the same for each language and that the results did not greatly differ. It should be noted that these are all veteran participants with much experience at CLEF.

<sup>5</sup> Ten groups submitted runs for monolingual English. We have included a graph showing the top 5 results but it must be remembered that the systems submitting these were actually focusing on the bilingual part of the task.

As usual in the CLEF monolingual task, the main emphasis in the experiments was on stemming and morphological analysis. The group from University of Neuchatel, which had the best overall performances for all languages, focused very much on stemming strategies, testing both light and aggressive stemmers for the Slavic languages (Bulgarian and Czech). For Hungarian they worked on decompounding. This group also compared performances obtained using word-based and 4-gram indexing strategies [9]. Another of the best performers, from Johns Hopkins University, normally uses an n-gram approach. Unfortunately, we have not received a paper from this group so cannot comment on their performance. The other group with very good performance for all languages was Opentext. In their working notes paper, this group also compared 4-gram results against results using stemming for all three languages. They found that while there could be large impacts on individual topics, there was little overall difference in average performance. In addition, their experiments confirmed past findings that indicate that blind relevance feedback can be detrimental to results, depending on the evaluation measures used [4]. The results of the statistical tests that can be found towards the end of our working notes paper show that the best results of these three groups did not differ significantly [8].

The group from Alicante also achieved good results testing query expansion techniques [10], while the group from Kolkata compared a statistical stemmer against other types of stemmers for Bulgarian, Czech and Hungarian with comparable results [11]. Czech is a morphologically complex language and the two Czech-only groups both used approaches involving morphological analysis and lemmatization [12], [13].

### 3.1 Results

Table 5 shows the top five groups for each target collection, ordered by mean average precision. The table reports: the short name of the participating group; the mean average precision achieved by the experiment; the DOI of the experiment; and the performance difference between the first and the last participant.

Figures comparing the performances of the top participants can be found in our working notes paper [8].

## 4 Main Stream Bilingual Experiments

The bilingual task was structured in three sub-tasks ( $X \rightarrow$  BG, CS, or HU target collection) plus a sub-task for non-European topic languages against an English target collection. A special sub-task testing Indian languages against the English collection was also organised in response to requests from a number of research groups working in India. For the bilingual to English task, participating groups also had to submit an English monolingual run, to be used both as baseline and also to reinforce the English pool. All groups participating in the Indian languages sub-task also had to submit at least one run in Hindi (mandatory) plus runs in other Indian languages (optional).

**Table 5.** Best entries for the monolingual track

Track	Rank	Participant	Experiment DOI	MAP
Bulgarian	1st	unine	10.2415/AH-MONO-BG-CLEF2007.UNINE.UNINEBG4	44.22%
	2nd	jhu-apl	10.2415/AH-MONO-BG-CLEF2007.JHU-APL.APLMOBGT4	36.57%
	3rd	opentext	10.2415/AH-MONO-BG-CLEF2007.OPENTEXT.OTBG07TDE	35.02%
	4th	alicante	10.2415/AH-MONO-BG-CLEF2007.ALICANTE.IRNBUEXP2N	29.81%
	5th	daedalus	10.2415/AH-MONO-BG-CLEF2007.DAEDALUS.BGFSBG2S	27.19%
	Difference			
Czech	1st	unine	10.2415/AH-MONO-CS-CLEF2007.UNINE.UNINECZ4	42.42%
	2nd	jhu-apl	10.2415/AH-MONO-CS-CLEF2007.JHU-APL.APLMOCSTD4	35.86%
	3rd	opentext	10.2415/AH-MONO-CS-CLEF2007.OPENTEXT.OTCS07TDE	34.84%
	4th	prague	10.2415/AH-MONO-CS-CLEF2007.PRAGUE.PRAGUE01	34.19%
	5th	daedalus	10.2415/AH-MONO-CS-CLEF2007.DAEDALUS.CSFSCS2S	32.03%
	Difference			
Hungarian	1st	unine	10.2415/AH-MONO-HU-CLEF2007.UNINE.UNINEHU4	47.73%
	2nd	opentext	10.2415/AH-MONO-HU-CLEF2007.OPENTEXT.OTHU07TDE	43.34%
	3rd	alicante	10.2415/AH-MONO-HU-CLEF2007.ALICANTE.IRNBUEXP2N	40.09%
	4th	jhu-apl	10.2415/AH-MONO-HU-CLEF2007.JHU-APL.APLMOHUTD5	39.91%
	5th	daedalus	10.2415/AH-MONO-HU-CLEF2007.DAEDALUS.HUFSHU2S	34.99%
	Difference			
English (only for Bilingual X2EN participants)	1st	bombay-ltrc	10.2415/AH-MONO-EN-CLEF2007.BOMBAY-LTRC.IITB_MONO_TITLE_DESC	44.02%
	2nd	jhu-apl	10.2415/AH-MONO-EN-CLEF2007.JHU-APL.APLMOENTD5	43.42%
	3rd	nottingham	10.2415/AH-MONO-EN-CLEF2007.NOTTINGHAM.MONONT	42.74%
	4th	depok	10.2415/AH-MONO-EN-CLEF2007.DEPOK.UIQDTMONO	40.57%
	5th	hyderabad	10.2415/AH-MONO-EN-CLEF2007.HYDERABAD.ENTD_QMENG07	40.16%
	Difference			

We were disappointed to only receive runs from one participant for the  $X \rightarrow$  BG, CS, or HU tasks. Furthermore, the results were quite poor; as this group normally achieves very good performance, we suspect that these runs were probably corrupted in some way. For this reason, we decided to disregard them as being of little significance. Therefore, in the rest of this section, we only comment on the  $X \rightarrow$  EN results.

We received runs using the following topic languages: Amharic, Chinese, Indonesian and Oromo plus, for the Indian sub-task, Bengali, Hindi, Marathi and Telugu<sup>6</sup>.

For many of these languages few processing tools or resources are available. It is thus very interesting to see what measures the participants adopted to overcome this problem. Here below, we briefly glance at some of the approaches and techniques adopted. For more details, see the papers cited.

The top performance in the bilingual task was obtained by an Indonesian group; they compared different translation techniques: machine translation using Internet resources, transitive translation using bilingual dictionaries and French and German as pivot languages, and lexicons derived from parallel corpus created by translating all the CLEF English documents into Indonesian using a commercial MT system. They found that they obtained best results using the MT system together with query expansion [14].

<sup>6</sup> Although topics had also been requested in Tamil, in the end they were not used.

The second placed group used Chinese for their queries and a dictionary based translation technique. The experiments of this group concentrated on developing new strategies to address two well-known CLIR problems: translation ambiguity, and coverage of the lexicon [15]. The work by [16] which used Amharic as the topic language also paid attention to the problems of sense disambiguation and out-of-vocabulary terms. They found that pseudo-relevance feedback improved their performance considerably.

The third performing group also used Indonesian as the topic language; unfortunately we have not received a paper from them so far so cannot comment on their approach. An interesting paper, although slightly out of the task as the topic language used was Hungarian was [17]. This group used a machine readable dictionary approach but also applied Wikipedia hyperlinks for query term disambiguation and exploited bilingual Wikipedia articles for dictionary extension. The group testing Oromo used linguistic and lexical resources developed at their institute; they adopted a bilingual dictionary approach and also tested the impact of a light stemmer for Oromo on their performance with positive results [18].

The groups using Indian topic languages tested different approaches. The group from Kolkata submitted runs for Bengali, Hindi and Telugu to English using a bilingual dictionary lookup approach [19]. They had the best performance using Telugu probably because they carried out some manual tasks during indexing. A group from Bangalore tested a statistical MT system trained on parallel aligned sentences and a language modelling based retrieval algorithm for a Hindi to English system [20]. The group from Bombay had the best overall performances; they used bilingual dictionaries for both Hindi and Marathi to English and applied term-to-term cooccurrence statistics for sense disambiguation [21]. The Hyderabad group attempted to build bilingual dictionaries using topical similarity by choosing vocabulary from a web search engine index and demonstrated that such dictionaries perform very well even with fewer entries [18]. Interesting work was done by the group from Kharagpur which submitted runs for Hindi and Bengali. They attempted to overcome the lack of resources for Bengali by using phoneme-based transliterations to generate equivalent English queries from Hindi and Bengali topics [23].

## 4.1 Results

Table 6 shows the best results for the bilingual task. The performance difference between the best and the fifth placed group is given (in terms of average precision). Again both pooled and not pooled runs are included in the best entries for each track, with the exception of Bilingual X  $\rightarrow$  EN.

For bilingual retrieval evaluation, a common method to evaluate performance is to compare results against monolingual baselines. This year we can only comment on the results for the bilingual to English tasks. The best results were obtained by a system using Indonesian as a topic language. This group achieved 88.10% of the best monolingual English IR system. This is a good result considering that Indonesian is not a language for which a lot of resources and machine-readable dictionaries are available. It is very close to the best results obtained

Table 6. Best entries for the bilingual task

Track	Rank	Part.	Lang.	Experiment DOI	MAP
Bulgarian	1st	jhu-apl	en	10.2415/AH-BILI-X2EG-CLEF2007.JHU-APL.APLBIENBGT4	7.33%
	Difference				
Czech	1st	jhu-apl	en	10.2415/AH-BILI-X2CS-CLEF2007.JHU-APL.APLBIENBST4	21.43%
	Difference				
English	1st	depok	id	10.2415/AH-BILI-X2EN-CLEF2007.DEPOK.UIQTDTOGGLEFB10D10T	38.78%
	2nd	nottingham	zh	10.2415/AH-BILI-X2EN-CLEF2007.NOTTINGHAM.GRAWOTD	34.56%
	3rd	jhu-apl	id	10.2415/AH-BILI-X2EN-CLEF2007.JHU-APL.APLBIIDENTDS	33.24%
	4th	hyderabad	om	10.2415/AH-BILI-X2EN-CLEF2007.HYDERABAD.OMTD07	29.91%
	5th	bombay-ltrc	hi	10.2415/AH-BILI-X2EN-CLEF2007.BOMBAY-LTRC.IITB_HINDI_TITLEDESC_DICE	29.52%
	Difference				
Hungarian	1st	jhu-apl	en	10.2415/AH-BILI-X2HU-CLEF2007.JHU-APL.APLBIENHUT5	29.63%
	Difference				

Table 7. Best entries for the bilingual Indian subtask

Track	Rank	Part.	Lang.	Experiment DOI	MAP
Hindi to English	1st	bombay-ltrc	hi	10.2415/AH-BILI-X2EN-CLEF2007.BOMBAY-LTRC.IITB_HINDI_TITLEDESC_DICE	29.52%
	2nd	msindia	hi	10.2415/AH-BILI-X2EN-CLEF2007.MSINDIA.2007_RBLM_ALL_CROSS_1000_POSSCORES	21.80%
	3rd	hyderabad	hi	10.2415/AH-BILI-X2EN-CLEF2007.HYDERABAD.HITD	15.60%
	4th	jadavpur	hi	10.2415/AH-BILI-X2EN-CLEF2007.JADAVPUR.AHBILIH2ENR1	10.86%
	5th	kharagpur	hi	10.2415/AH-BILI-X2EN-CLEF2007.KHARAGPUR.HINDITITLE	4.77%
	6th				
	Difference				
Bengali/ Hindi/ Marathi/ Telugu to English	1st	bombay-ltrc	hi	10.2415/AH-BILI-X2EN-CLEF2007.BOMBAY-LTRC.IITB_HINDI_TITLEDESC_DICE	29.52%
	2nd	msindia	hi	10.2415/AH-BILI-X2EN-CLEF2007.MSINDIA.2007_RBLM_ALL_CROSS_1000_POSSCORES	21.80%
	3rd	bombay-ltrc	mr	10.2415/AH-BILI-X2EN-CLEF2007.BOMBAY-LTRC.IITB_MAR_TITLE_DICE	21.63%
	4th	hyderabad	te	10.2415/AH-BILI-X2EN-CLEF2007.HYDERABAD.TEED	21.55%
	5th	jadavpur	te	10.2415/AH-BILI-X2EN-CLEF2007.JADAVPUR.AHBILITE2ENR1	11.28%
	6th	kharagpur	bn	10.2415/AH-BILI-X2EN-CLEF2007.KHARAGPUR.BENGALITITLEDESC	7.25%
	Difference				

last year for two well-established CLEF languages: French and Portuguese, when the equivalent figures were 93.82% and 90.91%, respectively.

## 4.2 Indian to English Subtask Results

Table 7 shows the best results for the Indian sub-task. The performance difference between the best and the last (up to 6) placed group is given (in terms of average precision). The first set of rows regards experiments for the mandatory topic language: Hindi; the second set of rows reports experiments where the source language is one of the other Indian languages.

It is interesting to note that in both sets of experiments, the best performing participant is the same. In the second set, we can note that for three (Hindi, Marathi, and Telugu) out of the four Indian languages used the performances of the top groups are quite similar.

The best performance for the Indian sub-task is 76.12% of the best bilingual English system (achieved by veteran CLEF participants) and 67.06% of the monolingual baseline, which is quite encouraging for a new task with languages where encoding issues and linguistic resources make the task difficult. This is in fact comparable with the performances of some newly introduced European languages. For example, we can compare them to those for Bulgarian and Hungarian in CLEF 2006:

- X → BG: 52.49% of best monolingual Bulgarian IR system;
- X → HU: 53.13% of best monolingual Hungarian IR system.

## 5 Robust Experiments

The robust task ran for the second time at CLEF 2007. It is an ad-hoc retrieval task based on test suites of previous CLEF campaigns. The evaluation approach is modified and a different perspective is taken. The robust task emphasizes the difficult topics by a non-linear integration of the results of individual topics into one result for a system [24,25]. By doing this, the evaluation results are interpreted in a more user-oriented manner. Failures and very low results for some topics hurt the user experience with a retrieval system. Consequently, any system should try to avoid these failures. This has turned out to be a hard task [26]. Robustness is a key issue for the transfer of research into applications. The robust task rewards systems which achieve a minimal performance level for all topics.

In order to do this, the robust task uses the geometric mean of the average precision for all topics (GMAP) instead of the mean average of all topics (MAP). This measure has also been used at a robust track at the Text Retrieval Conference (TREC) where robustness was explored for monolingual English retrieval [25]. At CLEF 2007, robustness was evaluated for monolingual and bilingual retrieval for three European languages.

The robust task at CLEF exploits data created for previous CLEF editions. Therefore, a test set with 100 topics can be used for the evaluation. Such a large number of topics allows a more reliable evaluation [27]. A secondary goal of the robust task is the definition of larger data sets for retrieval evaluation.

As described above, the CLEF 2007 robust task offered three languages often used in previous CLEF campaigns: English, French and Portuguese. The data used was developed during CLEF 2001 through 2006. Generally, the topics from CLEF 2001 until CLEF 2003 were used as training topics whereas the topics developed between 2004 and 2006 were the test topics on which the main evaluation measures are given.

Thus, the data used in the robust task in 2007 is different from the set defined for the robust task at CLEF 2006. The documents which need to be searched are articles from major newspapers and news providers in the three languages. Not all collections had been offered consistently for all CLEF campaigns, therefore, not all collections were integrated into the robust task. Most data from 1995 was omitted in order to provide a homogeneous collection. However, for Portuguese,

**Table 8.** Data for the Robust Task 2007

Language	Target Collection	Training Topic DOIs	Test Topic DOIs
English	LA Times 1994	10.2452/41-AH-10.2452/200-AH	10.2452/251-AH-10.2452/350-AH
French	Le Monde 1994 SDA 1994	10.2452/41-AH-10.2452/200-AH	10.2452/251-AH-10.2452/350-AH
Portuguese	Público 1995	—	10.2452/201-AH-10.2452/350-AH

**Table 9.** Best entries for the robust monolingual task

Track	Rank	Participant	Experiment DOI	MAP	GMAP
English	1st	reina	10.2415/AH-ROBUST-MONO-EM-TEST-CLEF2007.REINA.REINAENTDNT	38.97%	18.50%
	2nd	daedalus	10.2415/AH-ROBUST-MONO-EM-TEST-CLEF2007.DAEDALUS.EMFSEN22S	37.78%	17.72%
	3rd	hildesheim	10.2415/AH-ROBUST-MONO-EM-TEST-CLEF2007.HILDESHEIM.HIMDOENBRFNE	5.88%	0.32%
	4th				
	5th				
	<b>Difference</b>				<b>562.76%</b>
French	1st	unine	10.2415/AH-ROBUST-MONO-FR-TEST-CLEF2007.UNINE.UNINEFR1	42.13%	14.24%
	2nd	reina	10.2415/AH-ROBUST-MONO-FR-TEST-CLEF2007.REINA.REINAFRTDET	38.04%	12.17%
	3rd	jaen	10.2415/AH-ROBUST-MONO-FR-TEST-CLEF2007.JAEN.UJARTFR1	34.76%	10.69%
	4th	daedalus	10.2415/AH-ROBUST-MONO-FR-TEST-CLEF2007.DAEDALUS.FRFSFR22S	29.91%	7.43%
	5th	hildesheim	10.2415/AH-ROBUST-MONO-FR-TEST-CLEF2007.HILDESHEIM.HIMDOFRBRF2	27.31%	5.47%
	<b>Difference</b>				<b>54.27%</b>
Portuguese	1st	reina	10.2415/AH-ROBUST-MONO-PT-TEST-CLEF2007.REINA.REINAPTTDNT	41.40%	12.87%
	2nd	jaen	10.2415/AH-ROBUST-MONO-PT-TEST-CLEF2007.JAEN.UJARTPT1	24.74%	0.58%
	3rd	daedalus	10.2415/AH-ROBUST-MONO-PT-TEST-CLEF2007.DAEDALUS.PTFSPPT2S	23.75%	0.50%
	4th	xldb	10.2415/AH-ROBUST-MONO-PT-TEST-CLEF2007.XLDB.XLDBRBR16_10	1.21%	0.07%
	5th				
	<b>Difference</b>				<b>3,321.49%</b>

for which no training data was available, only data from 1995 was used. Table 8 shows the data for the robust task.

The robust task attracted 63 runs submitted by seven groups (CLEF 2006: 133 runs from eight groups). Effectiveness scores were calculated with version 8.0 of the `trec_eval` program which provides the *Mean Average Precision (MAP)*, while the DIRECT system version 2.0 was used to calculate the *Geometric Average Precision (GMAP)*. .

## 5.1 Robust Monolingual Results

Table 9 shows the best results for this task. The performance difference between the best and the last (up to 5) placed group is given (in terms of average precision). The results cannot be compared to the results of the CLEF 2005 and CLEF 2006 campaign in which the same topics were used because a smaller collection had to be searched. The working notes paper contains figures comparing the performances of the top participants for each language [8].

**Table 10.** Best entries for the robust bilingual task

Track	Rank	Participant	Experiment DOI	MAP	GMAP
French	1st	reina	10.2415/AH-ROBUST-BILI-X2FR-TEST-CLEF2007.REINA.REINAE2FTDNT	35.83%	12.28%
	2nd	unine	10.2415/AH-ROBUST-BILI-X2FR-TEST-CLEF2007.UNINE.UNINEBILFR1	33.50%	5.01%
	3rd	colesun	10.2415/AH-ROBUST-BILI-X2FR-TEST-CLEF2007.COLESUN.EN2FRTST4GRINTLGLU001	22.87%	3.57%
	4th				
	5th				
	Difference				54.27%

## 5.2 Robust Bilingual Results

Table 10 shows the best results for this task. The performance difference between the best and the last (up to 5) placed group is given (in terms of average precision). All the experiments were from English to French.

As previously stated for bilingual retrieval evaluation, a common method is to compare results against monolingual baselines. We have the following results for CLEF 2007:

- X → FR: 85.05% of best monolingual French IR system;

The results in Table 9 and Table 10 suggest that there is no difference between the rankings based on MAP and those based on GMAP. No position changes between system occur. However, a more thorough analysis of the CLEF 2006 robust results which included not only the best run of each system but also the other runs showed interesting results. The correlation between MAP and GMAP based rankings is decreasing with the number of topics for multilingual retrieval [32]. This result shows that the creation of larger test suites is necessary.

## 5.3 Approaches Applied to Robust Retrieval

The REINA system performed best in three tasks and seems to be well optimized for robust retrieval [34]. It applied different measures of robustness during the training phase in order to optimize the performance. A local query expansion technique added terms. Another system experimented with n-gram based translation for bi-lingual retrieval which requires no languages specific components [36]. SINAI tried to increase the robustness of the results by expanding the query with an external knowledge source [33]. This is a typical approach in order to obtain additional query terms and avoid zero hits in case of out of vocabulary problems. Contrary to standard query expansion techniques, the new terms form a second query and results of both initial and second query are integrated under a logistic fusion strategy. The Daedalus group submitted experiments with the Miracle system [35]; BM25 weighting without blind relevance feedback was applied. For detailed descriptions of all the robust experiments, see the Robust section in these Proceedings.



## 6 Conclusions

We have reported the results of the ad hoc cross-language textual document retrieval track at CLEF 2007. This track is considered to be central to CLEF as for many groups it is the first track in which they participate and provides them with an opportunity to test their text retrieval systems and compare performance between monolingual and cross-language runs, before perhaps moving on to more complex system development and subsequent evaluation. This year, the monolingual task focused on central European languages while the bilingual task included an activity for groups that wanted to use non-European topic languages and languages with few processing tools and resources. Each year, we also include a task aimed at examining particular aspects of cross-language text retrieval. Again this year, the focus was examining the impact of “hard” topics on performance in the “robust” task.

The paper also describes in some detail the creation of the pools used for relevance assessment this year and includes observations on their stability. We also performed a number of statistical tests on the results with the aim of determining what differences between runs appear to be real as opposed to differences that are due to sampling issues. Unfortunately, for reasons of space, we are unable to report the results here. The interested reader is again referred to our on-line working notes paper.

Although there was quite a good participation in the monolingual Bulgarian, Czech and Hungarian tasks and the experiments report some interesting work on stemming and morphological analysis, we were very disappointed by the lack of participation in bilingual tasks for these languages. On the other hand, the interest in the task for non-European topic languages was encouraging and the results reported can be considered positively.

The robust task has analyzed the performance of systems for older CLEF data under a new perspective. A larger data set which allows a more reliable comparative analysis of systems was assembled. Systems needed to avoid low performing topics. Their success was measured with the geometric mean (GMAP) which introduces a bias on poor performing topics. Results for the robust task for mono-lingual retrieval on English, French and Portuguese collections as well as for bi-lingual retrieval from English to French are reported. Robustness can also be interpreted as the fitness of a system under a variety of conditions. The definition on what robust retrieval means has to continue.

As a result of discussions at the workshop, the CLEF 2008 Ad Hoc track has been considerably revolutionized. We offer a totally new main task for monolingual and cross-language search on library catalogue records, organised in collaboration with The European Library (TEL). We also offer more traditional mono- and bilingual ad-hoc retrieval tasks on a Persian newspaper corpus: the Hamshahri collection. This is the first time that we offer a non-European target collection in CLEF. The 2008 “robust” task proposes monolingual and bilingual tasks on a word sense disambiguated (WSD) collection of news documents in English. The goal of the task is to test whether WSD can be used beneficially for retrieval systems. The results will be reported and discussed at the CLEF 2008 workshop.

## Acknowledgements

We should like to acknowledge the enormous contribution of the groups responsible for topic creation and relevance assessment. In particular, we thank the group responsible for the work on Bulgarian led by Kiril Simov and Petya Osenova, the group responsible for Czech led by Pavel Pecina<sup>7</sup>, and the group responsible for Hungarian led by Tamás Váradi and Gergely Bottyán. These groups worked very hard under great pressure in order to complete the heavy load of relevance assessments in time.

## References

1. Paskin, N. (ed.): The DOI Handbook – Edition 4.4.1. International DOI Foundation (IDF) (2006) [last visited 2007, August 30], <http://dx.doi.org/10.1000/186>
2. Braschler, M.: CLEF 2003 - Overview of results. In: Peters, C., Braschler, M., Gonzalo, J., Kluck, M. (eds.) CLEF 2003. LNCS, vol. 3237, pp. 44–63. Springer, Heidelberg (2004)
3. Tomlinson, S.: Sampling Precision to Depth 10000 at CLEF 2007. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 57–64. Springer, Heidelberg (2008)
4. Tomlinson, S.: Sampling Precision to Depth 10000: Evaluation Experiments at CLEF 2007. In: Nardi, A., Peters, C. (eds.) Working Notes for the CLEF 2007 Workshop (2007) [last visited May 2008], <http://www.clef-campaign.org/>
5. Braschler, M., Peters, C.: CLEF 2003 Methodology and Metrics. In: Peters, C., Gonzalo, J., Braschler, M., Kluck, M. (eds.) CLEF 2003. LNCS, vol. 3237, pp. 7–20. Springer, Heidelberg (2004)
6. Di Nunzio, G.M., Ferro, N.: Appendix A: Results of the Core Tracks – Ad-hoc Bilingual and Monolingual Tasks. In: Nardi, A., Peters, C. (eds.) Working Notes for the CLEF 2007 Workshop (2007) [last visited May 2008], <http://www.clef-campaign.org/>
7. Di Nunzio, G.M., Ferro, N.: Appendix B: Results of the Core Tracks – Ad-hoc Robust Bilingual and Monolingual Tasks. In: Nardi, A., Peters, C. (eds.) Working Notes for the CLEF 2007 Workshop (2007) [last visited May 2008], <http://www.clef-campaign.org/>
8. Di Nunzio, G.M., Ferro, N., Peters, C., Mandl, T.: CLEF 2007: Ad Hoc Track Overview. In: Nardi, A., Peters, C. (eds.) Working Notes for the CLEF 2007 Workshop [last visited May 2008], <http://www.clef-campaign.org/>
9. Dolamic, L., Savoy, J.: Stemming Approaches for East European Languages. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 37–44. Springer, Heidelberg (2008)
10. Noguera, E., Llopis, F.: Applying Query Expansion Techniques to Ad Hoc Monolingual Tasks with the IR-n system. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 45–48. Springer, Heidelberg (2008)
11. Majumder, P., Mitra, M., Pal, D.: Bulgarian, Hungarian and Czech Stemming using YASS. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 49–56. Springer, Heidelberg (2008)
12. Češka, P., Pecina, P.: Charles University at CLEF 2007 Ad-Hoc Track. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 33–36. Springer, Heidelberg (2008)

---

<sup>7</sup> Ministry of Education of the Czech Republic, project MSM 0021620838.

13. Ircing, P., Müller, L.: Czech Monolingual Information Retrieval Using Off-The-Shelf Components – the University of West Bohemia at CLEF 2007 Ad-Hoc track. In: Nardi, A., Peters, C. (eds.) Working Notes for the CLEF 2007 Workshop (2007) [last visited May 2008], <http://www.clef-campaign.org/>
14. Adriani, M., Hayurani, H., Sari, S.: Indonesian-English Transitive Translation for Cross-Language Information Retrieval. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 127–133. Springer, Heidelberg (2008)
15. Zhou, D., Truran, M., Brailsford, T.: Disambiguation and Unknown Term Translation in Cross Language Information Retrieval. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 64–71. Springer, Heidelberg (2008)
16. Argaw, A.A.: Amharic-English Information Retrieval with Pseudo Relevance Feedback. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 119–126. Springer, Heidelberg (2008)
17. Schönhofen, P., Benczúr, A., Bíró, I., Csalogány, K.: Cross-Language Retrieval with Wikipedia. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 72–79. Springer, Heidelberg (2008)
18. Pingali, P., Tune, K.K., Varma, V.: Improving Recall for Hindi, telugu, Oromo to English CLIR. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 103–110. Springer, Heidelberg (2008)
19. Bandyopadhyay, S., Mondal, T., Naskar, S.K., Ekbal, A., Haque, R., Godavarthy, S.R.: Bengali, Hindi and Telugu to English Ad-hoc Bilingual task at CLEF 2007. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 88–94. Springer, Heidelberg (2008)
20. Jagarlamudi, J., Kumaran, A.: Cross-Lingual Information Retrieval System for Indian Languages. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 80–87. Springer, Heidelberg (2008)
21. Chinnakotla, M.K., Ranadive, S., Damani, O.P., Bhattacharyya, P.: Hindi-English and Marathi-English Cross Language Information Retrieval Evaluation. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 111–118. Springer, Heidelberg (2008)
22. Pingali, P., Varma, V.: IIIT Hyderabad at CLEF 2007 – Adhoc Indian Language CLIR task. In: Nardi, A., Peters, C. (eds.) Working Notes for the CLEF 2007 Workshop [last visited May 2008] <http://www.clef-campaign.org/>
23. Mandal, D., Gupta, M., Dandapat, S., Banerjee, P., Sarkar, S.: Bengali and Hindi to English CLIR Evaluation. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 95–103. Springer, Heidelberg (2008)
24. Robertson, S.: On GMAP: and Other Transformations. In: Yu, P.S., Tsotras, V., Fox, E.A., Liu, C.B. (eds.) Proc. 15th International Conference on Information and Knowledge Management (CIKM 2006), pp. 78–83. ACM Press, New York (2006)
25. Voorhees, E.M.: The TREC Robust Retrieval Track. SIGIR Forum 39, 11–20 (2005)
26. Savoy, J.: Why do Successful Search Systems Fail for Some Topics. In: Cho, Y., Wan Koo, Y., Wainwright, R.L., Haddad, H.M., Shin, S.Y. (eds.) Proc. 2007 ACM Symposium on Applied Computing (SAC 2007), pp. 872–877. ACM Press, New York (2007)
27. Sanderson, M., Zobel, J.: Information Retrieval System Evaluation: Effort, Sensitivity, and Reliability. In: Baeza-Yates, R., Ziviani, N., Marchionini, G., Moffat, A., Tait, J. (eds.) Proc. 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2005), pp. 162–169. ACM Press, New York (2005)

28. Hull, D.: Using Statistical Testing in the Evaluation of Retrieval Experiments. In: Korfhage, R., Rasmussen, E., Willett, P. (eds.) Proc. 16th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 1993), pp. 329–338. ACM Press, New York (1993)
29. Conover, W.J.: Practical Nonparametric Statistics, 1st edn. John Wiley and Sons, New York (1971)
30. Judge, G.G., Hill, R.C., Griffiths, W.E., Lütkepohl, H., Lee, T.C.: Introduction to the Theory and Practice of Econometrics, 2nd edn. John Wiley and Sons, New York (1988)
31. Tague-Sutcliffe, J.: The Pragmatics of Information Retrieval Experimentation, Revisited. In: Spack Jones, K., Willett, P. (eds.) Readings in Information Retrieval, pp. 205–216. Morgan Kaufmann Publisher, Inc., San Francisco (1997)
32. Mandl, T., Womser-Hacker, C., Ferro, N., Di Nunzio, G.: How Robust are Multilingual Information Retrieval Systems? In: Proc. 2008 ACM SAC Symposium on Applied Computing (SAC), pp. 1132–1136. ACM Press, New York (2008)
33. Martínez-Santiago, F., Ráez, A.M., Garcia Cumbresas, M.A.: SINAI at CLEF Ad-Hoc Robust Track 2007: Applying Google Search Engine for Robust Cross-lingual Retrieval. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 137–142. Springer, Heidelberg (2008)
34. Zazo, A., Berrocal, J.L.A., Figuerola, C.G.: Improving Robustness Using Query Expansion. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 143–147. Springer, Heidelberg (2008)
35. González-Cristóbal, J.-C., Goñi-Menoyo, J.M., Goñi-Menoyo, J., Lana-Serrano, S.: MIRACLE Progress in Monolingual Information Retrieval at Ad-Hoc CLEF 2007. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 156–159. Springer, Heidelberg (2008)
36. Vilares, J., Oakes, M., Vilares Ferro, M.: English-to-French CLIR: A Knowledge-Light Approach through Character N- Grams Alignment. In: Peters, C., et al. (eds.) CLEF 2007. LNCS, vol. 5152, pp. 148–155. Springer, Heidelberg (2008)