

The Multiple Language Question Answering Track at CLEF 2003

Bernardo Magnini*, Simone Romagnoli*, Alessandro Vallin*
Jesús Herrera**, Anselmo Peñas**, Víctor Peinado**, Felisa Verdejo**
Maarten de Rijke***

- * ITC-irst, Centro per la Ricerca Scientifica e Tecnologica
Via Sommarive, 38050 Povo (TN), Italy.
{magnini,romagnoli,vallin}@itc.it
- ** UNED, Spanish Distance Learning University, Dpto. Lenguajes y Sistemas Informaticos
Ciudad Universitaria, c./Juan del Rosal 16, 28040 Madrid, Spain.
{jesus.herrera,anselmo,victor,felisa}@lsi.uned.es
- *** Language and Inference Technology Group, ILLC, University of Amsterdam
Nieuwe Achtergracht 166, 1018 WV Amsterdam, The Netherlands.
mdr@science.uva.nl

Abstract

This paper reports on the pilot question answering track that was carried out within the CLEF initiative this year. The track was divided into monolingual and bilingual tasks: monolingual systems were evaluated within the frame of three non-English European languages, Dutch, Italian and Spanish, while in the cross-language tasks an English document collection constituted the target corpus for Italian, Spanish, Dutch, French and German queries. Participants were given 200 questions for each task, and were allowed to submit up to two runs per task with up to three responses (either exact answers or 50 bytes long strings) per question.

We give here an overview of the track: we report on each task and discuss the creation of the multilingual test sets and the participants' results.

1 Introduction

The question answering (QA) track at TREC-8 represented the first attempt to emphasise the importance and foster research on systems that could extract relevant and precise information rather than documents. Question answering systems are designed to find answers to open domain questions in a large collection of documents. QA development has acquired an important role among the scientific community because it entails research in both natural language processing and information retrieval (IR), putting the two disciplines in contact. Differently from the IR scenario, a QA system processes questions formulated into natural language (instead of keyword-based queries) and retrieves answers (instead of documents).

The past TREC conferences laid the foundations for a formalised and widely accepted evaluation methodology of QA systems, but the three tracks organised so far focused just on monolingual systems for the English language, which constitutes a drawback we tried to address. We were mainly interested in testing multilingual systems, and in particular to push the QA community into designing them. As the number of the participants and the results achieved by their systems show, we can argue that in the field of multilingual QA there is much work to do. Within the frame of planning and co-ordinating the research on question answering, outlined in Maybury's roadmap, multilingual QA has a pivotal role and should deserve much attention in the next years. Multilinguality represents a new area in QA research, and a challenging issue toward the development of more complex systems [Maybury 2002].

Multilinguality enables the user to pose a query in a language that is different from the language of the reference corpus. The cross-language perspective could be quite useful when the required information is not available in the user's language (as it often happens surfing the web) and in particular it fits for the cultural situation in Europe, where different languages co-exist and are in contact, although English has become a widespread and standardised means of communication. In a multilingual environment, QA systems and other natural language processing resources could even contribute to conserve endangered languages that are progressively losing importance and prestige, in the effort to ensure their survival, as in the case of the 'Te Kaitito' bilingual question answering system for English and Ma_ri [Knott et al. 2001].

Our activity, and in particular the production of two multilingual test sets that constitute reusable resources, can be regarded as a valuable contribution to the development of such cross-language systems [Burger et al. 2001]. The evaluation of cross-language resources is the key issue of the CLEF initiative, so our question answering track could not be limited to the English language. On the contrary, we attempted to raise interest on other European languages, like Italian, Spanish, Dutch, German and French. The basic novelty in comparison with the

past TREC QA campaigns was the introduction of bilingual tasks, in which non-English queries are processed to find responses in an English document collection.

2 QA at CLEF

Our pilot question answering track was structured in both monolingual and bilingual tasks. We organised three monolingual tasks for Dutch, Italian and Spanish, in which the questions, the corpus and the responses were in the same language. In contrast, in the cross-language tasks we had Italian, Spanish, Dutch, French or German queries that searched for answers in an English document collection. In output, the systems had to retrieve English answers.

2.1 Monolingual Tasks

Unlike previous TREC QA tracks, we focused on the evaluation and on the production of reusable resources for non-English QA systems. The monolingual tasks were designed for three different languages: Dutch, Italian and Spanish. For each language we generated 200 queries, 180 of which were completely shared between all the three tasks. Participants were given the questions and the corresponding monolingual corpus: the task consisted in returning automatically, i.e. with no manual intervention, a ranked list of [docid, answer] pairs per question such that the retrieved document supported the answer. Participants were given 200 questions for each task, and were allowed to submit up to two runs per task with up to three responses per query. They could return either exact answers or 50 bytes long strings that contained the answer, although they were not allowed to use both modalities within the same run. Following the TREC model, we formulated 20 questions that had no known answer in the corpora: systems indicated their belief that there was no answer in the document collection by returning “NIL” instead of the [docid, answer] pair.

The monolingual Italian question answering task was planned and carried out under the co-ordination of the Italian Centro per la Ricerca Scientifica e Tecnologica (ITC-irst), that was in charge for the supervision of the whole QA track. We could use the document collections released at CLEF 2002, made up of articles drawn from a newspaper (*La Stampa*) and a press agency (*SDA*) of the year 1994. The entire Italian target corpus was 200 Mb wide (about 27 millions words) and it was made available to registered participants at the end of last January, so that they could test their systems using the document collection well in advance.

The UNED NLP group (Spanish Distance Learning University), as Spanish member of the CLEF consortium, was in charge for the monolingual Spanish task. The collection we were allowed to employ was the one released at CLEF 2002, i.e. more than 200,000 news from *EFE Press Agency* of the year 1994.

The Language and Inference Technology Group at the University of Amsterdam took care of the monolingual Dutch task. The collection used was the CLEF 2002 Dutch collection, which consists of two full years of the *Algemeen Dagblad* and *NRC Handelsblad* newspapers, adding up to about 200,000 documents of 540 Mb.

2.2 Cross-Language Tasks

Our interest in developing QA systems for languages other than English was not the only achievement we pointed at: the great novelty introduced in the CLEF QA track was multilinguality, whose potentialities are currently out of the scope of the TREC competition. Cross-language QA systems are crucially important when the language of the query and the language of the document collection are different, and in multicultural situations such a possibility is far from being remote. Searching information in the World Wide Web for instance is often difficult because the document retrieved is in a language we cannot understand. In this sense the cross-language tasks we organised represent a good chance to push the QA community to design and evaluate multilingual systems.

The cross-language tasks consisted in searching an English corpus to find English responses to queries posed in a different language. The target document collection we used was a corpus made up of *Los Angeles Times* articles of the year 1994, that was the same employed in last year’s CLEF campaign. We translated into five languages the original two hundred English questions we generated, so we were able to organise five different bilingual tasks: Dutch, French, German, Italian and Spanish. As in the monolingual tasks, participants had to process 200 questions (15 had no answer in the corpus) posed in one of the five languages and could choose to submit either exact answers or 50 bytes strings, without mixing them in the same run.

2.3 Participants

Eight groups took part in this pilot question answering track, and a total of seventeen runs were submitted, three using 50 bytes long strings as answers and the other fourteen, in compliance with last year’s TREC conditions, returning exact answers. The fact that most participants chose to retrieve exact answers shows that many have made the transition from more or less long strings to precise responses.

Figure 1 below shows the name of the participants, the task in which they participated and the filename of their runs. It is interesting to notice that all the participants except the DFKI group had already participated in some previous TREC QA campaigns.

GROUP	TASK	RUN NAME
DLSI-UA Uni. of Alicante, Spain	Monolingual Spanish	alicex031ms alicex032ms
UVA Uni. of Amsterdam, the Netherlands	Monolingual Dutch	uamsex031md uamsex031md
ITC-irst Italy	Monolingual Italian Bilingual Italian	irstex031mi irstst032mi irstex031bi irstex032bi
ISI Uni. of Southern California, USA	Bilingual Spanish	isixex031bs isixex032bs
/	Bilingual Dutch	/
DFKI Germany	Bilingual German	dfkist031bg
CS-CMU Carnegie Mellon Uni., USA	Bilingual French	lumox031bf lumox032bf
DLTG Uni. of Limerick, Ireland	Bilingual French	dltgex031bf dltgex032bf
RALI Uni. of Montreal, Canada	Bilingual French	udemst031bf udemex032bf

Figure 1: participants in the CLEF question answering track

Note that the fifth and sixth letters in the run names show whether the responses are exact answers (ex) or 50 bytes long strings (st).

Three teams took part in the monolingual tasks, submitting a total of six runs. We had only one participant in each language, which is quite disappointing because no comparison can be made between similar runs. Anyway, since the question set for all the monolingual tasks was the same (except the NIL questions), the monolingual runs can be compared to some extent. Four teams initially registered for the monolingual Italian task, but unfortunately only one, the ITC-irst group, actually participated. Similarly, only the University of Alicante took part in the monolingual Spanish task submitting two runs of exact answers, although three other groups expressed their intention of participation. As for the monolingual Dutch task, the University of Amsterdam with its two runs of exact answers was the only participant.

Six groups participated in the cross-language tasks, submitting eleven runs. The challenging novelty of the cross-language question answering attracted more participants than the monolingual tasks: the bilingual French task was chosen by three groups, while no one tested their system in the bilingual Dutch.

3 Test Sets

From a potential user’s point of view, a question answering system should be able to process natural language queries and return precise and unambiguous responses, drawn from a large reference corpus. Thus, in every evaluation campaign like the one we conducted, a set of well formulated questions is required. Since they should reflect real requests posed by humans, such questions must sound spontaneous and realistic. On the other hand, they must be clear, simple and factoid, i.e. related to facts, events, physical situations, so that the answers can be retrieved without inference. All the necessary information to answer the questions must be straightforwardly available and consequently included in the document collection searched by the systems. For this reason no external knowledge of the world should be required and the queries should deal with practical, concrete matters, rather than with abstract notions, that depend on personal opinion or reasoning.

The creation of the question sets for both the tasks – thoroughly described in the article “Creating the DISEQuA Corpus: a Test Set for Multilingual Question Answering” in these proceedings – entailed much work in terms of queries selection and answers verification. In order to establish some common criteria of comparison between the several languages involved, we decided to provide the participants, independently from the language, with the same queries. Thus, we created two collections of two hundred questions each, translated into different languages: one for the monolingual tasks and the other one for the cross-language tasks. As a result, we put together two reusable linguistic resources that can be useful for the QA community but also for other NLP fields, such as Machine Translation. The test set for the monolingual tasks in particular represents a multilingual collection of queries with their answers in different corpora.

3.1 Gold Standard for the Monolingual Tasks

The benchmark collection of queries and responses for the Dutch, Italian and Spanish monolingual tasks was the result of a joint effort between the co-ordinators, who decided to share the test sets in the three languages. Our activity can be roughly divided into four steps:

1. *Production of a pool of 200 candidate questions with their answers in each language.* These queries were formulated on the basis of the topics released by CLEF for the retrieval tasks of the year 2000, 2001 and/or 2002. The CLEF topics, i.e. a set of concepts chosen with the aim of covering the main events occurred in the years 1994 and/or 1995, allowed us to pose questions independently from the document collection. In this way we avoided any influence in the contents and in the formulation of the queries. Questions were posed according to common guidelines: they had to be generally short and fact-based, unrelated to subjective opinions. They could not ask for definitions (i.e. “Who is Bill Clinton”) and they had to have just one unique and unambiguous item as response, which means that we avoided questions asking for multiple items like those used in the TREC list task. Three groups of native speakers, one for each language, were involved in this work and searched the correct answers. A question has an answer in the reference corpus if a document contains the correct response without any inference implying knowledge outside the document itself.
2. *Selection of 150 questions from each monolingual set.* Since our aim was to build a test set of shared queries that would find answers in all the monolingual corpora, each group chose 150 questions from its candidate pool and translated them into English, thus a larger collection of 450 queries was put together. English constituted a sort of inter-language we used to shift from one language to another, but in this phase we were aware that there was the risk of changing unwarily the content of the questions during the translation. Each group chose its 150 questions taking into consideration that they would be processed by the other two, so the most general queries, that were likely to find a response in the other two corpora, were selected. Those that were too strictly related to the specific issues of a country were discarded.
3. *Processing of the shared questions.* Once we had collected a pool of 450 questions that had response in one of the corpora, we detected the duplicates and eliminated them. Quite surprisingly, we found thirteen couples of queries that had an identical meaning, although the formulation could be slightly different. Then each group translated back from English the 300 questions provided by the other co-ordinators and verified whether they had an answer in its corpus.
4. *Selection of the final 200 questions.* At this point, about 450 different questions had been formulated and translated into Dutch, English, Italian and Spanish. All of them had at least one answer in at least one language (other than English), and more than 200, obtained by merging the data of the second cross-verification, proved to have at least one answer in all the three monolingual document collections. Our goal was to provide the QA participants with 200 questions, including a small rate of NIL queries, i.e. questions that do not have any known answer in the corpus. We agreed that the 10% of the test set was a reasonable amount of NIL questions, that were first introduced in QA evaluation at TREC-10 (2001). So we selected 180 questions from those that had a response in all the three corpora, and each group completed its monolingual test set adding 20 NIL questions, that were necessarily different for each task. Taking into consideration seven general classes of questions, we tried to balance the final test set of 180 questions, that is composed of: 45 entries that ask for the name or role of a PERSON, 40 that pertain a LOCATION, 31 a MEASURE, 23 an ORGANISATION, 19 a DATE, 9 a concrete OBJECT, while 13, due to their vagueness, can be labelled with OTHER.

The result of the question development phase is a useful and reusable multilingual question set, whose entries are structured in a XML format, as shown in the example of Figure 2. More details are given in the paper “Creating the DISEQuA Corpus”.

```

<qa cnt="1" type="DATE">
  <language val="ITA" original="TRUE">
    <question assessor="ALE">
      Quando e' avvenuta la riunificazione delle due Germanie?
    </question>
    <answer n="1" idx="SDA19941115.00073">
      nel 1989
    </answer>
  </language>
  <language val="SPA" original="FALSE">
    <question assessor="Anselmo">
      ¿Cuándo se produjo la reunificación de Alemania?
    </question>
    <answer n="1" idx="EFE19941108-04388">
      1989
    </answer>
    <answer n="2" idx="EFE19941108-04508">
      1989
    </answer>
  </language>
  <language val="DUT" original="FALSE">
    <question assessor="LIT">
      Wanneer vond de Duitse hereniging plaats?
    </question>
    <answer n="1" idx="NH19940128-0161">
      in 1989
    </answer>
  </language>
  <language val="ENG" original="FALSE">
    <question assessor="">
      When did the reunification of East and West Germany take place?
    </question>
    <answer n="1" idx="-1">
      SEARCH[in 1989]
    </answer>
  </language>
</qa>1

```

Figure 2: gold standard format of a question for the monolingual tasks

3.2 Gold standard for the Cross-Language Tasks

While in the monolingual tasks we had three different document collections and three sets of questions, all the bilingual tasks had one English target corpus. For this reason we generated 200 English queries and verified manually that each of them (except 15 NIL) had at least an answer. Then the questions were translated into each language. As in the monolingual test sets, translators were asked to be as faithful as possible to the original English version, in fact we were aware that every translation could be different from the source.

Because of organisational problems encountered shortly before the test set creation deadline, three Italian native speakers at ITC-irst had to take on the job, even though there was a high risk of inconsistencies that may have affected the quality of the question set as a resource.

Due to time constraints we could not compile a large pool of general questions independently from the corpus and then verify them. Instead, we chose an alternative approach: we randomly selected a document from the collection (while trying to select news with a worldwide importance, avoiding sections that deal with local politics or issues too strictly related to Los Angeles counties) and picked up a text snippet that was relevant, long and interesting enough to get a question out of it. For instance, from the following passage

¹ In the files downloaded by participants, the questions for each monolingual QA task were written in another format, and numbered differently. The questions above appeared in the following way in the three monolingual test sets:

M DUT 0097 Wanneer vond de Duitse hereniging plaats?

M ITA 0077 Quando è avvenuta la riunificazione delle due Germanie?

M SPA 0013 ¿Cuándo se produjo la reunificación de Alemania?

The government has banned foods containing intestine or thymus from calves because a new scientific study suggested that they might be contaminated with the infectious agent of bovine spongiform encephalopathy, commonly called "mad cow disease".

we drew the question ‘What is another name for the “mad cow disease”?’.

Finally, we obtained a benchmark corpus in which each question appears in six languages, as the tag attribute <language val> in figure 3 shows:

```

<qa cnt="4" type="OTHER">
  <language val="ENG" original="TRUE">
    <question assessor="ALE">
      What is another name for the "mad cow disease"?
    </question>
    <answer n="1" idx="LA091194.0096">
      bovine spongiform encephalopathy
    </answer>
  </language>
  <language val="ITA" original="FALSE">
    <question assessor="ALE">
      Qual è un altro nome per la "malattia della mucca pazza"?
    </question>
    <answer n="1" idx="">
      SEARCH[bovine spongiform encephalopathy]
    </answer>
  </language>
  <language val="SPA" original="FALSE">
    <question assessor="">
      ¿Qué otro nombre recibe la enfermedad de las vacas locas?
    </question>
    <answer n="1" idx="">
      SEARCH[bovine spongiform encephalopathy]
    </answer>
  </language>
  <language val="DUT" original="FALSE">
    <question assessor="">
      Wat is een andere naam voor "gekke-koeienziekte"?
    </question>
    <answer n="1" idx="">
      SEARCH[bovine spongiform encephalopathy]
    </answer>
  </language>
  <language val="GER" original="FALSE">
    <question assessor="">
      Was ist ein anderer Name für "Rinderwahnsinn"?
    </question>
    <answer n="1" idx="">
      SEARCH[bovine spongiform encephalopathy]
    </answer>
  </language>
  <language val="FRE" original="FALSE">
    <question assessor="">
      Quel autre nom donne-t-on à la "maladie de la vache folle"?
    </question>
    <answer n="1" idx="">
      SEARCH[bovine spongiform encephalopathy]
    </answer>
  </language>
</qa>

```

Figure 3: gold standard format of a question for the bilingual tasks

4 Results

Participants had one week to process the questions. Since no manual intervention of any kind was allowed, we asked participants to freeze their systems before downloading the queries from our “QA @ CLEF” website.² Before the start of the evaluation exercise, we released detailed guidelines with the necessary information about the required format of the submissions. We also put online a checking routine with which participants could make sure that their responses were in compliance with that.

4.1 Response Format

Since we allowed to submit both exact answers and 50 bytes long strings, we could not evaluate these two formats together. For this reason, we divided our track into two subtasks with separated evaluations. The required format of the answers in both subtasks was the same, but we decided to draw up two separate results.

Figure 4 shows an example of a participant’s submissions, where the first column indicates the question number, provided by the organisers, and the string in the second one represents the unique identifier for a system and a run: the last two characters in this case show that the task is the bilingual Italian, and the fifth and sixth characters give information about the kind of responses retrieved in this run, i.e. exact answers.

The third field in the response format was the answer rank, which was crucially important for the evaluation of the system accuracy. Participants had to return the questions in the same order in which they had been downloaded, i.e. unranked. On the contrary, they had to rank their responses by confidence, putting in the first place the surest answer.

The integer or floating point score number of the fourth column justified the answer ranking. This field was not compulsory, and the systems that had no scoring strategies could set the value to default 0 (zero).

The docid, i.e. the unique identifier of the document that supports the given answer, is placed in the fifth column. If the system maintained that there was no answer in the corpus or if it could not find one, the docid was replaced by the string “NIL”.

The answer string had to be given in the last field of the response, that was left empty when the docid was substituted by “NIL”.

0001	irstex031bi	1	3253	LA011694-0094	Modern Art
0001	irstex031bi	2	1776	LA011694-0094	UCLA
0001	irstex031bi	3	1251	LA042294-0050	Cultural Center
0002	irstex031bi	1	9	NIL	
0003	irstex031bi	1	484	LA012594-0239	1991
0003	irstex031bi	2	106	LA012594-0239	Monday
0004	irstex031bi	1	154	LA072294-0071	Clark
0004	irstex031bi	2	117	LA072594-0055	Huber
0004	irstex031bi	3	110	LA072594-0055	Department

Figure 4: examples of responses drawn from the first bilingual run submitted by ITC-irst

4.2 Judgements and Evaluation Measures

Each single answer was judged by human assessors, who assigned to each response a unique label: either right, or wrong, or unsupported or inexact. Assessors were told to judge the submissions from a potential user’s point of view, because the evaluation should take into consideration the future portability of QA systems. They analysed both the answers themselves and the context, i.e. the document that supported the answer, in which they appeared.

Answers were judged to be incorrect (W) when the answer-string did not contain the answer or when the answer was not responsive. In contrast, a response was considered to be correct (R) when the answer-string consisted of nothing more than the exact, minimal answer (or contained the correct answer within the 50 bytes long string) and when the document returned supported the response. Unsupported answers (U) were correct but it was impossible to infer that they were responsive from the retrieved document. Answers were judged as non-exact (X) when the answer was correct and supported by the document, but the answer string missed bits of the response or contained more than just the exact answer.

In addition, we outlined some common criteria to distinguish and properly evaluate exact answers. We outlined general rules to apply in several cases: as regards the date of specific events that ended in the past, both day and year are normally required (unless the question refers only to the year), but if the day cannot be retrieved, the

² <http://clef-qa.itc.it>

year is normally sufficient. For instance, if a system answered the question “When did Napoleon die?” returning “5th May”, it would be judged as incorrect. On the other hand, both “May 5, 1821” and “1821” could be correct exact answers. Actually, no clear definitions of exact answer have been formalised, yet. Discussing the issue, we noticed that, generally speaking, articles and prepositions do not invalidate an “exact” answer. So, both “July, 9” and “on the 9th of July” are exact answers. Similarly, appositions should not represent a problem, as well. So for instance, “1957”, “year 1957” and “in the year 1957” should be exact answers, though someone could object that (with dates) “year” is redundant. When a query asks for a measure, the unit of measure can be accepted, too. So, both “30” and “30 degrees” are exact.

Concerning NIL answers, they are correct if neither human assessors nor systems have found any answer before or after the assessment process. If there is an answer in the collection, NIL is evaluated as incorrect. A NIL answer means that the system *believes* that there is not an answer for that question in the collection. There is no way for systems to explicitly indicate that they do not know or cannot find the answer for a question.

Question and judged responses	Comment
<p><i>What museum is directed by Henry Hopkins?</i></p> <p>W 1 irstex031bi 1 3253 LA011694-0094 Modern Art U 1 irstex031bi 2 1776 LA011694-0094 UCLA X 1 irstex031bi 3 1251 LA042294-0050 Cultural Center</p>	<p>The second answer was correct but the document retrieved was not relevant. The third response missed bits of the name, and was judged non-exact.</p>
<p><i>Where did the Purussaurus live before becoming extinct?</i></p> <p>W 2 irstex031bi 1 9 NIL</p>	<p>The system erroneously “believed” that the query had no answer in the corpus, or could not find one.</p>
<p><i>When did Shapour Bakhtiar die?</i></p> <p>R 3 irstex031bi 1 484 LA012594-0239 1991 W 3 irstex031bi 2 106 LA012594-0239 Monday</p>	<p>In the questions that asked for the date of an event, the year was often regarded as sufficient.</p>
<p><i>Who is John J. Famalaro accused of having killed?</i></p> <p>W 4 irstex031bi 1 154 LA072294-0071 Clark R 4 irstex031bi 2 117 LA072594-0055 Huber W 4 irstex031bi 3 110 LA072594-0055 Department</p>	<p>The second answer, that returned the victim’s last name, was considered sufficient and correct, since in the document retrieved no other people named “Huber” were mentioned.</p>

Figure 5: examples of judged responses drawn from the first bilingual run submitted by ITC-irst

In strict evaluation, only correct answers (R) scored points, while in lenient evaluation the unsupported responses (U) were considered to be correct, too. The score of each question was the reciprocal of the rank for the first answer to be judged correct, which means that each query could receive either 1, or 0, or 0.333, or 0.5 points, depending on the confidence ranking.

The basic evaluation measure was the Mean Reciprocal Rank (MRR), that represents the mean score over all questions. MRR takes into consideration both recall and precision of the systems’ performance, and can range between 0 (no correct responses) and 1 (all the 200 queries have a correct answer at position one). Figures 6 and 7 below summarise the QA track results and show that the systems achieved better results in the monolingual than in the bilingual tasks, where the drop in performance is possibly due to the cross-language step. The QA system developed by ITC-irst proved to be the most accurate among those that participated, and the mean reciprocal rank scored in the monolingual Italian using 50 bytes long strings as answers was the highest of the whole QA track.

Answer responsiveness and exactness were in the opinion of human assessors, whose judgement could be different, as in everyday life we have different criteria to determine whether a response is good or not. During the evaluation of most of the runs, two different assessors judged each single question (each question of the bilingual runs were judged by three NIST assessors) and in case of discrepancies, they discussed their opinion and tried to reach an agreement. Whenever they could not agree, another person took the final decision.

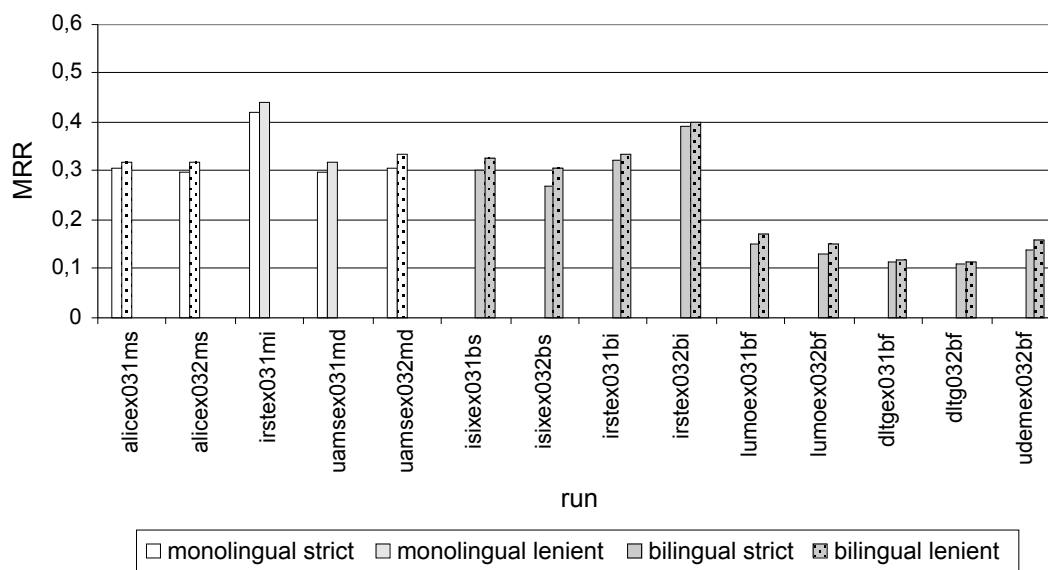
After the submission deadline had passed, we detected some mistakes in the questions. In particular, a blunder persisted in the Italian queries: we wrongly put an apostrophe after the contraction of the question word “quale” (“which”/“what”). We found 21 cases in the monolingual test set and 17 cases in the bilingual one. In the TREC campaigns the questions that contain mistakes are excluded from the evaluation, but, considering that the form “qual’e’era” is quite common in Italian and that a QA system should be robust enough to recognise variant spellings, we decided to keep those queries. For the sake of completeness, we calculated precision and recall without the questions with that mistake, and we obtained just a very minor variation of the values (around 1%).

Translation could be the source of mistakes, as well. In the monolingual Spanish questions collection, “minister of Foreign Affairs” was erroneously translated as “president of Foreign Affairs” during the question sharing between the Italian and the Spanish co-ordinators.

EXACT ANSWERS RUNS

	GROUP	TASK	RUN NAME	MRR		No. of Q. with at least one right answer		NIL Questions	
				strict	lenient	strict	lenient	returned	correctly returned
MONO-LINGUAL TASKS	DLSI-UA	monolingual Spanish	alicex031ms	.307	.320	80	87	21	5
			alicex032ms	.296	.317	70	77	21	5
	ITC-irst	monolingual Italian	irstex031mi	.422	.442	97	101	4	2
	UVA	monolingual Dutch	uamsex031md	.298	.317	78	82	200	17
uamsex032md			.305	.335	82	89	200	17	
CROSS-LANGUAGE TASKS	ISI	bilingual Spanish	isixex031bs	.302	.328	69	77	4	0
			isixex032bs	.271	.307	68	78	4	0
	ITC-irst	bilingual Italian	irstex031bi	.322	.334	77	81	49	6
			irstex032bi	.393	.400	90	92	28	5
	CS-CMU	bilingual French	lumox031bf	.153	.170	38	42	92	8
			lumox032bf	.131	.149	31	35	91	7
	DLTG	bilingual French	dltgex031bf	.115	.120	23	24	119	10
			dltgex032bf	.110	.115	22	23	119	10
	RALI	bilingual French	udemex032bf	.140	.160	38	42	3	1

Figure 6: summary statistics of the exact answer runs

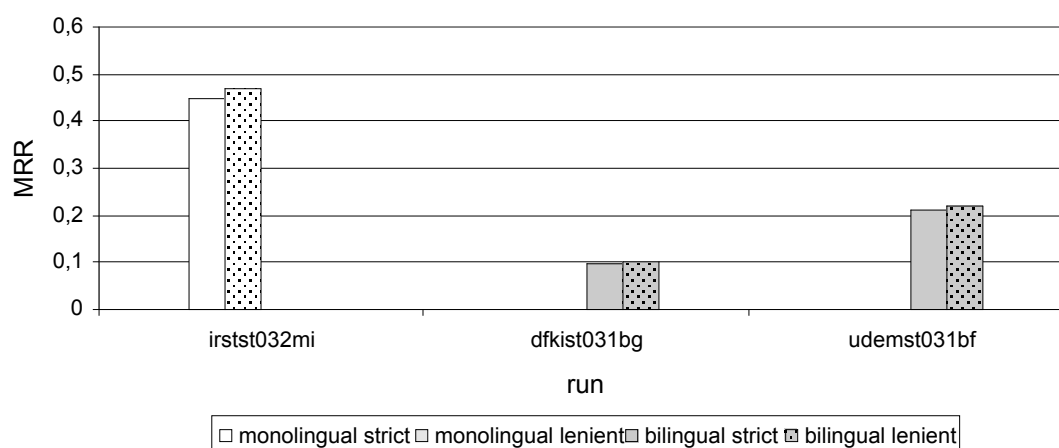


Graph 1: Mean Reciprocal Rank score of the exact answer runs

50 BYTES LONG ANSWER RUNS

	GROUP	TASK	RUN NAME	MRR		No. Q with at least one Right answer		NIL Questions	
				strict	lenient	strict	lenient	returned	correctly returned
MONO-LINGUAL TASKS	ITC-irst	Monolingual Italian	irstst032mi	.449	.471	99	104	5	2
CROSS-LANGUAGE TASKS	DFKI	Bilingual German	dfkist031bg	.098	.103	29	30	18	0
	RALI	Bilingual French	udemst031bf	.213	.220	56	58	4	1

Figure 7: summary statistics of the 50 bytes answer runs



Graph 2: Mean Reciprocal Rank score of the 50 bytes long answer runs

As can be noticed in graphs 1 and 2, strict and lenient evaluation results do not differ much. More strikingly, the performance of the cross-language systems turned out to be quite low, which suggests that multilinguality is a field that requires much more attention and investigation.

5 Conclusions and Future Perspectives

The first European evaluation of non-English QA systems has given rise to useful resources for future multilingual QA developments. It has allowed us to establish and test a methodology and criteria for both the test suit production and the assessment procedure. Unfortunately, the CLEF QA Track did not receive the expected attention in terms of participation, and in most tasks just one group submitted its results. Actually, twelve research groups registered and were interested into participating, but some of them could not adjust their system on time. This suggests that the debate and the activities on multilingual QA have a certain appeal on the community, even though much challenging work remains to be done. We can be pleased of the outcome of this

pilot QA evaluation exercise, and we hope that the results and the resources we developed will encourage many other groups to participate in future campaigns.

Cross-linguality has always been out of the scope of the TREC QA tracks, and our pilot QA at CLEF hopefully represents a first step in the direction of more sophisticated evaluation campaigns of multilingual systems. In our track, we provided five non-English question sets but just one English target document collection: in the future we could have several reference corpora in different languages, many different question sets and answers translated into different languages. Multilinguality provides us with the opportunity to experiment with different approaches, exploring many potential applications: for instance, we could think about developing intelligent systems that taking into consideration the language and the text coverage, select the most useful target corpus to search the answer for a particular question posed in a particular language. The possibilities are manifold, and our cross-language tasks can be considered just a starting point.

6 Acknowledgements

The work described in this paper has been supported by the Autonomous Province of Trento, in the framework of the WebFAQ Project, by the Spanish Government (MCyT, TIC-2002-10597-E) and by the Netherlands Organization for Scientific Research (NWO) under project numbers 612-13-001, 365-20-005, 612.069.006, 612.000.106, 220-80-001, and 612.000.207.

The authors wish to acknowledge the contribution and support given by Franca Rossi, Elisabetta Fauri, Pamela Forner and Manuela Speranza at ITC-irst, who helped us in the generation and verification of the questions for both the monolingual and bilingual tasks.

Stephan Busemann at DFKI and Jian-Yun Nie at the University of Montreal took on the job of translating into German and French the questions for the bilingual tasks, and their contribution was fundamental.

We wish to thank the NIST, and in particular Ellen Voorhees and Donna Harman, for judging all the bilingual runs and for providing us with the necessary resources and feedback we needed to organise this pilot track.

We are also grateful to Carol Peters (ISTI-CNR) and Charles Callaway (ITC-irst) who, as English native speakers, edited the questions for the cross-language tasks.

We wish to thank Henry Chinaski, Vera Hollink, and Valentin Jijkoun for their help in the development and assessment of the monolingual Dutch task.

References

[Braschler and Peters 2002]

Braschler, M. and Peters, C.: "The CLEF Campaigns: Evaluation of Cross-Language Information Retrieval Systems", in UPGRADE (The European Online Magazine for the IT Professional), Vol. III, Issue no. 3, (June 2002), pp. 78-81, CEPIS.

Electronic version available at <http://www.upgrade-cepis.org/issues/2002/3/up3-3Braschler.pdf>.

[Burger et al. 2001]

Burger, J., Cardie, C., Chaudhri, V., Gaizauskas, R., Harabagiu, S., Israel, D., Jacquemin, C., Lin, C.-Y., Maiorano, S., Miller, G., Moldovan, D., Ogden, B., Prager, J., Riloff, E., Singhal, A., Shrihari, R., Strzalkowski, T., Voorhees, E., Weishedel, R.: "Issues Tasks and Program Structures to Roadmap Research in Question & Answering (Q&A)", (2001).

Electronic version available at http://www-nlpir.nist.gov/projects/duc/papers/qa.Roadmap-paper_v2.doc.

[CLEF 2003 QA Track Guidelines]

"CLEF 2003 Question Answering Track: Guidelines for the Monolingual and Bilingual Tasks".

Electronic version available at <http://clef-qa.itc.it/guidelines.htm>.

[Knott et al. 2001]

Knott A., Bayard I., de Jager S., Smith L., Moorfield J. and O'Keefe R.: "A question-answering system for English and Ma_ri", in Proceedings of the Fifth Biannual Conference on Artificial Neural Networks and Expert Systems (ANNES), University of Otago, November 2001, pp. 223--228.

[Liddy 2002]

Liddy, E.D.: "Why are People Asking these Questions? A Call for Bringing *Situation* into Question-Answering System Evaluation", in LREC Workshop Proceedings on Question Answering – Strategy and Resources, Grand Canary Island, Spain, 2002.

[Magnini et al. 2001]

Magnini B., Negri M., Prevete R., Tanev H.: "Multilingual Question/Answering: the DIOGENE System", in Proceedings of the Tenth Text REtrieval Conference (TREC 2001), NIST Special Publication 500-250, 2001.

[Magnini 2002]

Magnini, B.: "Evaluation of Cross-Language Question Answering Systems", proposal presentation held at the CLEF Workshop 2002.

Electronic Version available at: <http://clef.iei.pi.cnr.it:2002/workshop2002/presentations/q-a.pdf>.

[Maybury 2002]

Maybury, M.: "Toward a Question Answering Roadmap", 2002.

Electronic version available at:

www.mitre.org/work/tech_papers/tech_papers_02/maybury_toward/maybury_toward_qa.pdf

[Voorhees 1999]

Voorhees, E. M.: "The TREC-8 Question Answering Track Report", in Proceedings of the Eighth Text REtrieval Conference (TREC-8), NIST Special Publication 500-246, 1999.

[Voorhees 2001]

Voorhees, E. M.: "Overview of the TREC 2001 Question Answering Track", in Proceedings of the Tenth Text REtrieval Conference (TREC 2001), NIST Special Publication 500-250, 2001.

[Voorhees 2002]

Voorhees, E. M.: "Overview of the TREC 2002 Question Answering Track", in Proceedings of the Eleventh Text REtrieval Conference (TREC 2002), NIST Special Publication 500-251, 2002.

[Voorhees and Tice 2000]

Voorhees, E. M., Tice, D. M.: "Building a Question Answering Test Collection", in Proceedings of SIGIR2000. Athens, Greece, ACM Press, 2000.