

University of Wolverhampton at CLEF 2007

Georgiana Puşcaşu and Constantin Orăsan
Research Group in Computational Linguistics
University of Wolverhampton, UK
georgie@wlv.ac.uk and C.Orasan@wlv.ac.uk

Abstract

This paper reports on the participation of the University of Wolverhampton in the Multiple Language Question Answering (QA@CLEF) track of the CLEF 2007 campaign. We approached the Romanian to English cross-lingual task with a Question Answering (QA) system that processes a question in the source language (i.e. Romanian), translates the identified keywords into the target language (i.e. English), and finally searches for answers in the English document collection. We submitted one run of our system that has achieved an overall accuracy of 14%. Besides the difficulties posed by developing a monolingual QA system, the bottleneck in building a cross-lingual one is the lack of a reliable translation methodology from the source into the target language.

Categories and Subject Descriptors

H.3 [Information Storage and Retrieval]: H.3.1 Content Analysis and Indexing; H.3.3 Information Search and Retrieval; H.3.4 Systems and Software; I.2 [Artificial Intelligence]: I.2.7 Natural Language Processing

General Terms

Measurement, Performance, Experimentation

Keywords

Question Answering, Cross-lingual Question Answering, Natural Language Processing

1 Introduction

Question Answering (QA) [7] is defined as the task of providing an exact answer to a question formulated in natural language. Cross-lingual QA capabilities enable systems to retrieve the answer in one language (the target language) to a question posed in a different language (the source language).

Last year, a new Romanian-to-English (RO-EN) cross-lingual QA task was organised for the first time within the context of the CLEF campaign [10]. The task consisted in retrieving answers to Romanian questions in an English document collection. Four types of questions were considered: factoid, definition, list and temporally restricted (see [10] for a detailed description of each question type). This year's task was organised in a similar manner, with the exception that all questions were clustered in classes related to same topic, some of which even contain anaphoric references to other questions from the same topic class or to their answers. Besides the usual news collections employed in the search for answers, this year's novelty was the fact that Wikipedia articles could

also be used as answer source, which significantly increased the search space, making the task more difficult.

This is the first time a Romanian-English cross-lingual QA system fully developed at the University of Wolverhampton has participated in the QA@CLEF competition. This system adheres to the classical architecture of QA systems which includes three stages: question processing, information retrieval and answer extraction [7]. In addition, the cross-lingual capabilities are provided by a Romanian-to-English term translation module. This paper describes the development stages and evaluation results of our system. The rest of the paper is organised as follows: Section 2 provides an overall description of the system, while Sections 3, 4, 5 and 6 present the four embedded modules - question processing, term translation, passage retrieval and answer extraction respectively. Section 7 captures the evaluation results and their analysis. Finally, in Section 8, conclusions are drawn and future directions of system development are considered.

2 System overview

Question Answering systems normally adhere to a pipeline architecture consisting of three main stages: question analysis, passage retrieval and answer extraction [7]. For cross-lingual systems, the language barrier is usually crossed by employing free online translation services for translating the question from the source language into the target language [8, 14]. The QA process is then entirely performed in the target language by a monolingual QA system. There are also cross-lingual systems that automatically translate the document collection in the source language and then perform monolingual QA in the source language [2]. Another alternative approach involves monolingual QA in the source language and then translating the answer, but this approach is feasible only when document collections covering the same material are available both in the source and target languages [1].

Since we could not identify neither reliable translation services from Romanian into English for translating complete questions, nor English-Romanian full document translation tools, the first two approaches could not be adopted. In the case of the third approach, the impediment was the lack of a Romanian document collection equivalent to the English one. Therefore we adopted a slightly different approach where the question analysis is performed in the original source language without any translation in order to overcome the negative effect of full question translation on the overall accuracy of the system. Afterwards, in order to link the two languages involved in the cross-lingual QA setting, term translation is performed by means of bilingual resources and linguistic rules. The search for passages and answers is then performed in the target language documents using modules designed for that particular language. This approach has been previously adopted by Sutcliffe et al. [13] and Tanev et al. [15].

The architecture of our system consists of a four-module pipeline, where each module is responsible for a different stage in answering a question. These four modules are:

- 1) ***Question Processing Module***

This module receives as input a question in Romanian, parses it with a statistical part-of-speech (POS) tagger and with a shallow parser, and then uses this linguistic information to identify the type of the question and of the expected answer, the question focus, as well as the relevant keywords.

- 2) ***Term Translation Module***

This module is responsible for identifying all translation equivalents of each term identified in the question. The translation equivalents are generated by consulting bilingual resources and then assembled into terms in the target language by means of linguistic rules.

- 3) ***Passage Retrieval Module***

At this stage candidate snippets of text are retrieved from the English document collection on the basis of a query that includes the translation equivalents of all terms identified in the question.

4) *Answer Extraction Module*

This module, on the basis of the information extracted by the Question Processor, processes the snippets of text retrieved at the previous stage and identifies candidate answers restricted to the expected answer type. Then one answer is selected after ranking the resulting list of candidate answers.

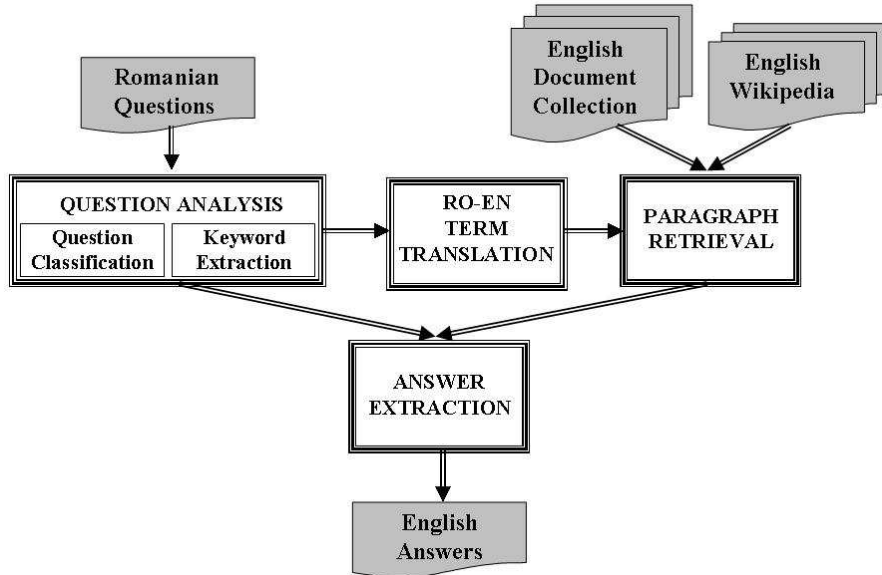


Figure 1: System Architecture and Functionality

Figure 1 illustrates the system architecture and functionality. The following four sections will present in more detail the functionality of each module.

3 Question Processing

This stage is mainly concerned with the identification of the semantic type of the entity sought by the question (expected answer type). In addition, it also provides the question focus, the question type and the set of keywords relevant for the question. To achieve these goals, our question analyser performs the following steps:

a) *POS-tagging, NP-chunking, Named Entity (NE) Extraction, Temporal Expression Identification*

The questions are first morpho-syntactically pre-processed using the TnT statistical part-of-speech tagger [3] trained on Romanian [17]. On the basis of this morpho-syntactic annotation, a rule-based shallow noun phrase (NP) chunker was implemented. A rule-based NE recogniser identifies the NEs which appear in the questions. Temporal expressions (TEs) are also detected using a Romanian TE identifier and normalizer based on the one previously developed for English by Puscasu [11].

b) *Question Focus Identification*

The question focus is the word or word sequence that defines or disambiguates the question, in the sense that it pinpoints what the question is searching for or what it is about. The question focus is considered to be either the noun determined by the question stem (for example in the question *What city hosted the Olympic Games in 2000?*, the focus is *city*) or the head noun of the first question NP if this NP comes before the question's main verb or if it follows the

verb “to be” (for example in the question *Who is the inventor of the polygraph?*, the focus is *inventor*).

c) ***Distinguishing the Expected Answer Type***

At this stage the category of the entity expected as an answer to the analysed question is identified. Our system’s answer type taxonomy distinguishes the following classes: PERSON, LOCATION, ORGANIZATION, TEMPORAL, NUMERIC, DEFINITION and GENERIC, and it was derived on the basis of questions asked in previous CLEF campaigns. The assignment of a class to an analysed question is performed using the question stem and the type of the question focus. The question focus type is detected using WordNet [5] sub-hierarchies specific to the categories PERSON / LOCATION / ORGANIZATION. We employ a pre-defined correspondence between each category and the ILI (InterLingual Index) codes of WordNet root nodes heading category-specific noun sub-trees. These ILI codes guide the extraction of category specific noun lists from the Romanian WordNet [18, 19]. In the case of ambiguous question stems (e.g. *What*), the resulted lists are searched for the head of the question focus, and the expected answer type is identified with the category of the corresponding list (for example, in the case of the question *In which country was Swann born?*, the question focus is *country*, noun found in the LOCATION list, therefore the associated expected answer type is LOCATION).

d) ***Inferring the Question Type***

This year, the QA@CLEF main task distinguishes among four question types: *factoid*, *definition*, *list* and *temporally restricted* questions¹. As temporal restrictions can constrain any type of question, we proceed by first detecting whether the question has the type *factoid*, *definition* or *list* and then test the existence of temporal restrictions. The question type is identified using two simple rules: for questions which ask for definitions of concepts, the assigned question type is *definition*; if the question focus is a plural noun, then the question type is *list*, otherwise the consider the question to be *factoid*. The temporal restrictions are identified using several patterns and the information provided by the TE identifier.

e) ***Keyword Set Generation***

The set of keywords is automatically generated by listing the question terms in decreasing order of their relevance, as follows: the question focus, the identified NEs and TEs, the remaining noun phrases, and all the non-auxiliary verbs present in the question. Given the grouping of questions into topics and the presence of anaphoric expressions pointing to terms situated in other questions belonging to the same topic, a shallow anaphora resolution mechanism was employed to expand the set of question keywords with other possibly relevant terms as described below. The expanded set of keywords is then passed on to the Term Translation module, in order to obtain English keywords for passage retrieval.

f) ***Resolution of anaphoric expressions***

One novelty introduced in this year’s competition was that questions were organised in clusters of related questions. In a number of cases, the links between questions were realised using anaphoric pronouns which meant that in order to obtain a more complete list of keywords, anaphora resolution was necessary. Given the difficulty of anaphora resolution it was not possible to employ a fully fledged anaphora resolution system. Instead, the set of keywords related to a question was expanded with the list of named entities present in the cluster. This was done for two reasons. On the one hand, investigation of the question clusters revealed that pronouns quite often refer to named entities in the cluster. On the other hand, given that the questions are related, it is possible that named entities present in the questions also co-occur in the same document. As a result, it is more likely to extract relevant documents with this expanded query. A number of questions referred to the result of the previous question. Currently, we took no steps to address this problem due to the fact that in our present system there is no way to feed the answer to a question back into the system.

¹For more details please refer to the track guidelines available at <http://clef-qa.itc.it/2007/guidelines.html>

4 Term Translation

At this stage two processes are carried out: term translation and query generation. The keywords extracted at the question processing stage are first translated with an approach similar to the one we employed last year when we participated together with two Romanian research groups in the same task at CLEF 2006 [12]. It does also resemble the one employed by Ferrandez et al. [6] within the same CLEF campaign, but in the English to Spanish cross-lingual task. After the process of term translation has finished, a query is generated by making a conjunction of all keywords. Each keyword is represented by all its translation equivalents grouped using the disjunction operator.

Term translation is achieved by employing WordNet and more specifically the ILI alignment between the English WordNet and all other WordNets developed as part of the EuroWordNet and BalkaNet projects. The underlying idea is that, given a Romanian word, the Romanian WordNet and its alignment to the English one, we identify all possible translations of the word by finding all the synsets containing it and crossing through the ILI alignment to the English side where the equivalent synsets are found. If the word to be translated does not appear in the Romanian WordNet, as is quite frequently the case, we search for it in other available dictionaries and preserve the first three translations. If still no translation is found, the word itself is considered as translation, an approach which works reasonably well for named entities.

In the case of multi-word terms, like most of the question noun phrases (NP), each NP word is translated individually using the method described above. After that, rules are employed to convert the Romanian syntax to English syntax, and to obtain the translation equivalents of a given term.

One drawback of this term translation method is that it proposes too many translations for a word due to the fact that it does not employ any word sense disambiguation. In order to address this problem, we implemented a ranking method which relies on information from parallel English-Romanian Wikipedia pages related to the question to be answered, but not necessary containing the actual answer. The assumption of this method was that the two sets of pages will contain more or less the same information, so it will be possible to find the most likely translation for the noun-verb pairs present in the question. Unfortunately, preliminary experiments revealed that the inclusion of this approach lead to the retrieval of a very small number of passages, many of which did not contain the answer to the question. Due to the time restrictions with this task, we were unable to properly tune the method to improve the quality of the passage retrieval module, and for this reason we did not employ it in this year's submission.

5 Passage Retrieval

The purpose of the passage retrieval module is to extract a list of passages from the document collection which may contain the answer to the question asked. This year's document collection consists of three distinct collections: English Wikipedia pages collected in November 2006, Los Angeles Times from 1994 and Glasgow Herald from 1995. This is the first time that Wikipedia has been included in the document collection and, as a result of the fact that it is several orders of magnitude bigger than the other two collections, the search space was significantly larger than in previous years, making the task more difficult. Given that the documents in each collection are formatted in different ways, each had to be indexed individually and processed in a slightly different manner. For indexing and retrieval, we used Lucene [9], an open source information retrieval library appropriate for local document collections and intranets.

The query proposed by the term translation module, including all possible translations of the question keywords, was used as starting point in extracting passages. In the initial experiments we tried to limit the number of translations used for each original keyword, but as a result, the number of retrieved snippets was too low. This can be explained by the fact that no disambiguation was performed and therefore it was possible that some of the translations were highly ranked and therefore included in the query, even though they were not appropriate. As explained before, attempts to order the keyword translations according to the likeliness of them being the correct

translation of the keyword did not lead to satisfactory results and therefore it was not used in this year's submission. In light of this, we decided to use all the translations identified for a keyword and linked them using the OR operator provided by Lucene.

We indexed the document collection in order to retrieve documents which contain the keywords, and not actual passages. We decided to take this approach because it offers more flexibility and allows better control of the methods which retrieve candidate passages. However, the drawback of this approach is that it needs to process each document individually and extract relevant passages. For this year's system we decided to retrieve only sentences. In order to do this, each sentence from the retrieved documents was scored on the basis of how many keywords, temporal expressions and named entities they contained. At present, up to 25 sentences with the highest scores are retrieved from each document, provided that their score is higher than a predefined threshold. This set of sentences is fed into the next module, the answer extractor.

6 Answer Extraction

Once candidate answer-bearing document passages (in our case sentences) have been selected, the answer extraction module starts with a merging of all passages retrieved for questions belonging to a certain topic. All retrieved passages are morpho-syntactically analysed and annotated with functional dependency information by employing Conexor's FDG Parser [16]. They are also parsed with the Named Entity Identifier embedded in the GATE (General Architecture for Text Engineering) toolkit [4], which recognises and classifies multi- or one-word strings as names of companies, persons, locations, etc.

Afterwards a question-based passage ranking is applied to the merged set of passages retrieved in response to queries derived from all topic specific questions. This set is ranked by using information that refers to the presence of the question focus, presence of question NEs, as well as of NEs belonging to the unified topic NE set, presence of other question elements (noun phrases and verb phrases), and presence of temporal and numeric expressions pertaining to the question.

The answer extraction process then addresses each type of expected answer type in a different manner, as follows:

- a) ***Expected answer type is a Named Entity such as PERSON, LOCATION, ORGANIZATION or MISCELLANEOUS (any other type of named entity)***

When the expected answer type is either a Named Entity or a NUMERIC / TEMPORAL entity, a text unit should be identified in the retrieved passages whose semantic type matches that of the expected answer. Named entities having the desired answer type are identified in the retrieved passages and added to the set of candidate answers. For each candidate answer, another score is computed on the basis of the passage score, the distance to other keywords and its frequency in the set of candidate answers. The candidate answer featuring the highest score is presented as the final answer. In the case of no candidate answer being found in the retrieved passages, the system returns NIL.

- b) ***Expected answer type is NUMERIC***

In the case of NUMERIC answers, there are several sub-categories we consider in our search for an answer: MONEY, PERCENTAGE, MEASURE and NUMERIC-QUANTITY (any other type of NUMERIC entity). Various patterns are defined for exact candidate answer identification, patterns that take into consideration either the format of certain numeric expressions or the presence of the question focus in the neighbourhood of a numeric expression. The process of ranking candidate answers relies on the same parameters as in the case of the Named Entity answer type.

- c) ***Expected answer type is TEMPORAL (i.e. a Temporal Expression)***

The sub-categories of TEMPORAL entities that guide the answer extraction process are: MILLENNIUM, CENTURY, DECADE, YEAR, MONTH, DATE, TIME, DURATION (this category also applies to questions asking about age) and FREQUENCY. Patterns have been

defined to extract from a certain temporal expression only that part having the required granularity (e.g. extracting from a temporal expression of granularity DATE like *25th of January 1993* only the YEAR, that is *1993*).

d) ***Expected answer type is GENERIC***

When the expected answer type is neither a Named Entity, nor a NUMERIC or TEMPORAL entity, the question focus is essential in finding the answer. The candidate answers are constrained to be hyponyms of the question focus head.

e) ***Expected answer type is DEFINITION***

When the expected answer is the definition of a concept, the processing is done in a different manner. Instead of using the passage extractor described in the previous section, it was decided to use a simpler approach. Wikipedia defines a large number of concepts, and therefore it was decided to first try to obtain the definition from the Wikipedia page associated to the concept. To this end, Lucene was used to return Wikipedia pages which contain the words from the concept to be defined in their title. Because this approach returned more than one document, a scoring method was implemented in order to rank the retrieved documents. If the document title contained words from the concept to be defined, the score of the document was boosted. In the case of words from the title not present in the concept, the score was penalised. Once the documents were ranked, regular expressions such as **X [is|are|was|were] [a|an|the] [possible definition]** were used to locate the answer to a question. A common problem with the documents extracted from Wikipedia is that quite often they do not have any real content and they are redirections to other pages which contain the real description of the concept. This problem had to be addressed before documents were scored. Whenever no answer could be located in Wikipedia, passages were extracted from the other two document collections using the passage retrieval module described in Section 5 and the regular expressions were then applied to them. Unfortunately, this fall-back approach performed quite poorly.

7 Evaluation Results

This section describes the results we obtained in our CLEF-2007 participation. We submitted only one run for the Romanian to English cross-lingual QA task. The methodology we employed targeted precision at the cost of recall, therefore we always chose to provide NIL answers for those questions we could not reliably locate a candidate answer in the retrieved passages. Apart from this, we have never returned more than one answer per question, but only the first ranked answer, when this could be identified.

Table 1 illustrates the detailed results achieved by our system. It is to be mentioned that our system is able to recognise at the Question Processing stage questions asking for LISTS, but the answer extractor does not tackle this type of questions.

	FACTOID	LIST	DEFINITION	TEMPORALLY RESTRICTED
RIGHT	15	0	13	0
WRONG	140	9	17	2
UNSUPPORTED	4	0	0	1
INEXACT	2	0	0	0
TOTAL	161	9	30	3
ACCURACY	9.32%	0.00%	43.33%	0.00%

Table 1: Detailed evaluation results

The overall accuracy of our cross-lingual QA system was evaluated at a generic score over all questions of 14%. An analysis of our system output revealed the fact that our system was unable to locate an answer and thus returned the answer NIL for 117 questions. It retrieved 83 answers, out of which 28 were correct, 49 were wrong, 4 unsupported and 2 inexact.

Unsupported answers are correct answers to a question, but the judge who evaluated the run considered the passage returned as a source for the answer not relevant enough for that particular question. Given that at this moment we do not have access to the correct answers and the expected support passages, it is difficult to judge whether the four retrieved passages are appropriate or not. For example, in the case of the Romanian question

Ce tip de animal a incercat Victor Bernal sa cumpere pe 25 ianuarie 1993? (which translates into English as *What kind of animal did Victor Bernal try to buy on the 25th of January 1993?*),

our returned answer was *gorilla* extracted from the following support passage:

The sting took place on Jan. 25, 1993, when Bernal and the others were escorted onto a DC-3 cargo plane parked in a remote corner of a small Miami airport to see the gorilla, crated for shipment.

which seems correct and justified by the presence of both Bernal's name and the date mentioned in the question, as well as the presence of the noun *gorilla*, which is a type of *animal*.

In the case of inexact answers, the answer-string contains a correct answer and the provided text-snippet supports it, but the answer-string is incomplete/truncated or is longer than the minimum amount of information required. For example, given the Romanian question

Ce meserie are Michael Barrymore? (which translates into English as *What is the occupation of Michael Barrymore?*),

our answer, evaluated as inexact, was *troubled comic* and the passage supporting it was:

Troubled comic Michael Barrymore last night received an ovation as his show, Strike It Lucky, was named Quiz Programme of the Year at the National Television Awards.

These errors can be corrected by improving the answer extractor with more specific rules as to the extent of the required answer.

A preliminary analysis of the incorrect and NIL answers showed that their main cause was the poor translation of the question keywords, this yielding either irrelevant or no passages being retrieved from the English document collection.

8 Conclusions

This paper described the development stages of our cross-lingual Romanian to English QA system, as well as our participation in the QA@CLEF campaign. Adhering to the generic QA system architecture, our system implements the three essential stages (question processing, passage retrieval and answer extraction), as well as a term translation module which provides cross-lingual capabilities by translating question terms from Romanian into English. It should be pointed out that this year our emphasis was less on fine tuning the system, and more on exploring the issues posed by the task and developing a complete system that can participate in the competition. Therefore, all four modules are still in a preliminary stage of development.

Our participation in the QA@CLEF campaign included only one run for the Romanian to English cross-lingual QA task. Our cross-lingual QA system achieved an overall accuracy of 14%. An in-depth analysis of our results at different stages in the QA process has revealed a number of future system improvement directions. The term translation module has a crucial influence over the performance of our system, and therefore will receive most of our attention. Apart from this, we will further investigate the ranking method for translation equivalents which relies on information from parallel English-Romanian Wikipedia pages in order to improve its performance, as we believe it is a promising research direction. We also intend to improve our answer extraction module by identifying a better answer ranking strategy.

9 Acknowledgements

The work presented in this paper has been partially supported by the EU funded project QALLME (FP6 IST-033860).

References

- [1] Johan Bos and Malvina Nissim. Cross-Lingual Question Answering by Answer Translation. In *Working Notes for the Cross Language Evaluation Forum (CLEF) 2006 Workshop*, Alicante, Spain, 2006.
- [2] Mitchell Bowden, Marian Olteanu, Pasin Suriyentrakorn, Jonathan Clark, and Dan Moldovan. LCC's PowerAnswer at QA@CLEF 2006. In *Working Notes for the Cross Language Evaluation Forum (CLEF) 2006 Workshop*, Alicante, Spain, 2006.
- [3] Thorsten Brants. TnT - a statistical part-of-speech tagger. In *Proceedings of the Sixth Conference on Applied Natural Language Processing (ANLP-2000)*, Seattle, WA, 2000.
- [4] Hamish Cunningham, Diana Maynard, Kalina Bontcheva, and Valentin Tablan. GATE: A framework and graphical development environment for robust NLP tools and applications. In *Proceedings of the 40th Anniversary Meeting of the Association for Computational Linguistics*, 2002.
- [5] Christiane Fellbaum, editor. *WordNet: An Electronic Lexical Database*. The MIT Press, 1998.
- [6] Sergio Ferrandez, Pilar Lopez-Moreno, Sandra Roger, Antonio Ferrandez, Jesus Peral, Xavier Alvarado, Elisa Noguera, and Fernando Llopis. AliQAn and BRILI QA Systems at CLEF 2006. In *Working Notes for the Cross Language Evaluation Forum (CLEF) 2006 Workshop*, Alicante, Spain, 2006.
- [7] Sanda Harabagiu and Dan Moldovan. Question Answering. In Ruslan Mitkov, editor, *Oxford Handbook of Computational Linguistics*, chapter 31, pages 560 – 582. Oxford University Press, 2003.
- [8] Valentin Jijkoun, Gilad Mishne, Maarten de Rijke, Stefan Schlobach, David Ahn, and Karin Muller. The University of Amsterdam at QA@CLEF2004. In *Working Notes for the Cross Language Evaluation Forum (CLEF) 2004 Workshop*, Bath, UK, 2004.
- [9] LUCENE. <http://lucene.apache.org/java/docs/>.
- [10] Bernardo Magnini, Danilo Giampiccolo, Pamela Forner, Christelle Ayache, Petya Osenova, Anselmo Peas, Valentin Jijkoun, Bogdan Sacaleanu, Paulo Rocha, and Richard Sutcliffe. Overview of the CLEF 2006 Multilingual Question Answering Track. In *Working Notes for the Cross Language Evaluation Forum (CLEF) 2006 Workshop*, Alicante, Spain, 2006.
- [11] Georgiana Puscasu. A Framework for Temporal Resolution. In *Proceedings of the 4th Conference on Language Resources and Evaluation (LREC2004)*, 2004.
- [12] Georgiana Puscasu, Adrian Iftene, Ionut Pistol, Diana Trandabat, Dan Tufis, Alin Ceausu, Dan Stefanescu, Radu Ion, Constantin Orasan, Iustin Dornescu, Alex Moruz, and Dan Cristea. Cross-Lingual Romanian to English Question Answering at CLEF 2006. In *Working Notes for the Cross Language Evaluation Forum (CLEF) 2006 Workshop*, Alicante, Spain, 2006.
- [13] Richard Sutcliffe, Michael Mulcahy, Igal Gabbay, Aoife O’Gorman, Kieran White, and Darina Slattery. Cross-Language French-English Question Answering using the DLT System at CLEF 2005. In *Working Notes for the Cross Language Evaluation Forum (CLEF) 2005 Workshop*, Vienna, Austria, 2005.

- [14] Hristo Tanev, Milen Kouylekov, Bernardo Magnini, Matteo Negri, and Kiril Ivanov Simov. Exploiting Linguistic Indices and Syntactic Structures for Multilingual Question Answering: ITC-first at CLEF 2005. In *Working Notes for the Cross Language Evaluation Forum (CLEF) 2005 Workshop*, Vienna, Austria, 2005.
- [15] Hristo Tanev, Matteo Negri, Bernardo Magnini, and Milen Kouylekov. The DIOGENE question answering system at CLEF-2004. In *Working Notes for the Cross Language Evaluation Forum (CLEF) 2004 Workshop*, Bath, UK, 2004.
- [16] Pasi Tapanainen and Timo Jaervinen. A Non-Projective Dependency Parser. In *Proceedings of the 5th Conference of Applied Natural Language Processing, ACL*, 1997.
- [17] Dan Tufis. Using a Large Set of EAGLES-compliant Morpho-Syntactic Descriptors as a Tagset for Probabilistic Tagging. In *Proceedings of the Second International Conference on Language Resources and Evaluation*, pages 1105 – 1112, Athens, Greece, May 2000.
- [18] Dan Tufis, Dan Cristea, and Sofia Stamou. BalkaNet: Aims, Methods, Results and Perspectives. A General Overview. In D. Tufis, editor, *Romanian Journal on Information Science and Technology. Special Issue on BalkaNet*. Romanian Academy, 2004.
- [19] Dan Tufis, Verginica Barbu Mititelu, Alexandru Ceausu, Luigi Bozianu, Catalin Mihaila, and Magda Manu. New developments of the Romanian WordNet. In *Proceedings of the Workshop on Resources and Tools for Romanian NLP*. 2006.