Blind Relevance Feedback and Named Entity based Query Expansion for Geographic Retrieval at GeoCLEF 2006

Kerstin Bischoff, Thomas Mandl, Christa Womser-Hacker

Information Science, University of Hildesheim,
Marienburger Platz 22
D-31141 Hildesheim, Germany

{mandl,womser}@uni-hildesheim.de

Abstract

The participation of the University of Hildesheim focused on the monolingual German and English and the bilingual German ←→ English tasks of GeoCLEF 2006. Based on the results of GeoCLEF 2005 the weighting and expansion of geographic named entities (NE) within a Boolean retrieval approach were combined. Since the best results 2005 were achieved with Blind Relevance Feedback (BRF), in which NEs seemed to play a crucial role, the effects of adding particular geographic NEs within the BRF are also explored. The paper presents a description of the system design, the submitted runs and preliminary results. A first analysis especially of inofficial post experiments indicates that geographic NEs can improve BRF and supports prior findings that the geographical expansion within a Boolean retrieval approach does not necessarily lead to better results – as is often assumed. The discussion is followed by planned future research activies.

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

General Terms

Measurement, Performance, Experimentation

Keywords

Cross-Language Information Retrieval, Evaluation, Geographic Information Retrieval Systems

1 Introduction

Many queries posted to web search engines are of geographic nature, i.e. the information searched for is restricted to a certain geographic region or place. Whereas the development of Geographic Information Systems (GIS) for using structured spatial data e.g. in urban or route planning and their enhancement for new location based services has long attracted much interest, Geographic Information Retrieval (GIR) is a comparatively new field of scientific research. Having participated in CLEF ad-hoc, GIRT, Web and ODQA tracks before the University of Hildesheim joins this year's GeoCLEF efforts to evaluate different approaches for improving the retrieval of geo-referenced information from unstructured data like newspaper articles (or web documents).

Given that in GeoCLEF 2005, where various approaches ranging from basic IR techniques to elaborated spatial indexing and retrieval methods were used [6], the most successful runs by Gey and Petras [7] were based on a fine-tuned BRF, we focused on the role of geographic NEs within the process of BRF. Gey and Petras [7] found that improvement through BRF seemed highly related to adding mostly proper names as "good" terms to the original query, thus the expansion by primarily geographic named entities (NE) should further improve retrieval quality.

In contrast, the results of last year's GeoCLEF track showed worse retrieval performance after manual or automatic expansion of geographic NEs to include their finer-grained sub regions resp. more specific location names; even combined with a Boolean approach the results were mixed [5,7,9]. Therefore the task for 2006 was especially designed to explore the usefulness of additional geographic information (mostly country names) provided in the narratives of the topics. Hence, we experimented with the recognition, weighting and expansion of such geographic NEs (not) using Boolean conjunction for German and English monolingual and German ←→ English bilingual retrieval.

2 Geographic Retrieval System

The system we augmented for this experimentation with (geographic) NEs in GIR is based on a retrieval system applied to ad-hoc retrieval in previous CLEF campaigns [8]. It uses Lucene's[1] technology for indexing and searching based both on similarity as well as the Boolean model. For bilingual runs topics were first translated by combining the translation tools Babelfish[2], Linguatec[3], FreeTranslation[4]. By merging multiple translation services the influence of wrong translations should be minimized as well as possible synonyms in the target language added to the query. Lexical analysis of queries resp. documents was realised using the University of Neuchatel's stopword lists[5] for English and German with added common CLEF words for topic formulation. Morphological analysis was done by the Lucene stemmer for German and the Snowball stemmer for English.

For Named Entity Recognition (NER) we employed the open source machine learning tool LingPipe[6], which identifies named entities and classifies them into the categories Person, Organization, Location and Miscellaneous according to a trained statistical model. For English we used the model provided by LingPipe, which is trained on newspaper data, and for German a model trained on an annotated corpus of German newspaper articles (Frankfurter Rundschau) [10] was used, because LingPipe has no predefined German model. NER was applied for query processing for weighting NEs and to generate Boolean queries of the type *concept AND geographical reference*. NER was also applied to the document collection for building an index with separate fields for the different categories of NEs to later allow for their systematic weighting within the BRF.

According to Gey and Petras [7] NEs seemed to be a crucial factor for successful BRF, which lead to above-average results for the respective runs at GeoCLEF 2005. Since the "good" terms to be added to the original query were mostly proper nouns, the addition of primarily geographic names during the BRF process should further improve retrieval quality. This might be particularly promising for imprecise regions like *Northern Germany* or other geographic names that are not to be found in a geographic thesaurus (*gazetteer*). Comprehensive gazetteers are often not publicly available. Even if so, adequate heuristics need to be established on which information to extract in order to effectively expand a query. Thus looking for frequent co-occurrences of the geographic NE of a query with other geographic NEs within the top-ranked documents may provide hints to the appropriate kind of information and enable the inclusion of more specific location names without the help of a gazetteer (e.g. *Ruhrgebiet – Bochum, Gelsenkirchen*).

The single steps in query processing and retrieval were carried out in our system in the following sequence (optional steps in parentheses):

- topic
- (translation)
- (NER and weighting)
- stopword removal
- stemming
- Boolean or ranked (Boolean) query
- (BRF with or without NE weighting; expansion Boolean or ranked)

3 Submitted Runs

After experimentation with the GeoCLEF data of 2005 we submitted runs differing in parameters and query processing steps. For monolingual English the base runs had no additional steps, i.e. no NER and weighting and no BRF. As mandatory, one run used only the title and description fields of the topics. The comparative run with

[1] http://lucene.apache.org/java/
[2] http://babelfish.altavista.com
[3] http://www.linguatec.de/onlineservices/pt
[4] http://www.freetranslation.com
[5] http://www.unine.ch/info/clef/
[6] http://www.alias-i.com/lingpipe

constant options also used the additional information from the narrative. To test the idea of geographic expansion through BRF we moreover submitted two runs (varying in respect of the used fields), in which (geographic) NEs were recognized and weighted as well as primarily extracted from the top-ranked 5 documents in the BRF process and added to the geographic clause of a Boolean query. In a fith run we did not combine these geographic entities from the BRF with the Boolean approach, but searched similarity based with this GeoBRF-expansion by 20 terms of the 5 best documents. In the monolingual German task we submitted the same two runs as for English, with NER and weighting, geographic BRF (GeoBRF) in form of 25 terms from 5 documents and with Boolean conjunction of all geographic NEs with the other terms. For comparison we had two runs with "traditional" BRF not highlighting any NEs at all. Since German base runs without any BRF performed poorly with respect to the topics of GeoCLEF 2005, we did not submit such an official run.

For the bilingual tasks German → English and English → German the run options, i.e. parameters and processing steps, were the same as in the runs for monolingual retrieval in the respective target language. Run descriptions and results measured as Mean Average Precision (MAP) are shown in Table 1 for monolingual runs and in Table 2 for bilingual runs.

Table 1. Results monolingual runs

Run Identifier	Language	Fields	NEs	BRF	Query	MAP
HIGeoenenrun1	English	TD	–		OR	16.76
HIGeoenenrun1n	English	TDN	–		OR	17.47
HIGeoenenrun2	English	TD	weighted	5 docs, 25 terms, GeoNEs and NEs weighted	AND	11.66
HIGeoenenrun2n	English	TDN	weighted	5 docs, 25 terms, GeoNEs and NEs weighted	AND	12.13
HIGeoenenrun3	English	TD	–	5 docs, 20 terms, GeoNEs weighted	OR	18.75
HIGeodederun4	German	TD	–	5 docs, 25 terms	OR	15.58
HIGeodederun4n	German	TDN	–	5 docs, 25 terms	OR	16.01
HIGeodederun6	German	TD	weighted	5 docs, 25 terms, GeoNEs and NEs weighted	AND	12.14
HIGeodederun6n	German	TDN	weighted	5 docs, 25 terms, GeoNEs and NEs weighted	AND	11.34

Table 2. Results bilingual runs

Run Identifier	Language	Fields	NEs	BRF	Query	MAP
HIGeodeenrun11	De → En	TD	–		OR	15.04
HIGeodeenrun11n	De → En	TDN	–		OR	19.03
HIGeodeenrun13	De → En	TD	weighted	5 docs, 25 terms, GeoNEs and NEs weighted	AND	14.56
HIGeodeenrun13n	De → En	TDN	weighted	5 docs, 25 terms, GeoNEs and NEs weighted	AND	15.65
HIGeodeenrun12	De → En	TD	–	5 docs, 20 terms, GeoNEs weighted	OR	16.03
HIGeodeenrun21	En → De	TD	–	5 docs, 25 terms	OR	11.86
HIGeodeenrun21n	En → De	TDN	–	5 docs, 25 terms	OR	13.15
HIGeodeenrun22	En → De	TD	weighted	5 docs, 25 terms, GeoNEs and NEs weighted	AND	09.69
HIGeodeenrun22n	En → De	TDN	weighted	5 docs, 25 terms, GeoNEs and NEs weighted	AND	10.46

Our results reached the average performance of all participants for monolingual and bilingual tasks. A first look at the MAPs for single topics revealed that performance was best for five or six topics, which were perhaps the more ad-hoc style ones (GC 29, 30, 32, 46, 48). The relatively low scores may indicate that geographic IR is a special task and outline the importance of evaluating methods to improve retrieval performance. Though geographic BRF adding 25 terms from 5 documents combined with the Boolean approach worsened performance for all tasks compared to the respective base runs, geographic BRF alone seemed to improve retrieval quality at least slightly. Query expansion by additional geographic information from the narrative field led neither to substantial improvement nor deterioration of performance. In fact, the only notable difference was found in bilingual German → English base runs, where improvement was mainly due to a higher accuracy for topic GC 38 (100% with using the narrative), a topic which only had one relevant document in the (pooled) collection.

Bilingual retrieval performance relies on translation quality. Measuring translation quality is difficult, but merging the three translation services seems to work well considering particularly wrong translations. Especially translating English into German resulted in many errors, while the opposite direction did produce much better translations: most often sentences were grammatically and stylistically incorrect – particularly if gerunds were involved in the English original –, but important keywords were correct. Choosing the wrong word from the dictionary (ambiguity in the target language) e.g. *banks of European rivers* → *Bänken von europäischen*

Flüssen; Länderparlamente → *Zustandsparlamente; forest fires* → *Waldfeuer* or confusion between parts of speech e.g. *car bombings near Madrid* → *Autombomben nähern sich Madrid* should have a more significant impact on retrieval quality. Problems with decompounding in translating geographic entities may be the source of e.g. *new narrow country* instead of *New England* or the *chew case mountains* instead of *Caucasus Mountains*. Missed translation or transliteration and the word by word translation of compound geographic NEs like *Middle East/Naher Osten; Vila Real* → *Vila actually; Oman* → *grandmas* can also be found.

Though it would be ideal to have a comprehensive gazetteer including name variants for all possible languages, merging could almost always assure that at least one correct translation was contributed to the automatically translated query. No single service seemed superior, but often different mistakes were made. Critical for all were only the *Westjordanland* (*Westbank), Berg-Karabach/ Nagorno-Karabakh*, the *Sea of Japan/ das Japanische Meer, Mediterranean sea (*→*Mittelmeermeer*) and *Ruhr area*. While for the last NE the opposite translation from German *Ruhrgebiet* to English produced valid variants, synonyms for concept terms could rarely be found. We will further analyse, if the improvement in the Boolean bilingual runs German → English compared to the respective monolingual runs can be attributed to the introduction of synonyms.

Combining multiple translation services could on the other hand considerably improve NER for German. The evaluation of the NER revealed an unsatisfying rate of approximately 45% correct recognition and classification of geographic entities (incl. *Warsaw Pact* and the *Eastern Bloc*) for German monolingual runs. After the translation the rate improved by 7 % (52 %). High improvement can be seen within the narrative of topic GC 37 (*Middle East*) where much better NER results were achieved. Listings of geographic entities within the narratives was a major difficulty for NER in German except for the last topic GC 50. Since the formulation of some topics did not vary between title and description as far as grammatical embedding of geographic NEs is concerned (e.g. GC 31, 33), all of them where missed. For English, an effect of translation was not observed. The comparatively much better NER rate decreased from 72% to 70%. Of course, (wrong and grammatically ill formed) translation also introduced some false positives, which would later be weighted high in the query. In general automatic translations of very low quality did not lead to false positives, but the NER module did the same mistakes or rather missed the same geographic entities. Especially difficult was the recognition of the correct boundaries of compound names like the *Sea of Japan, the East Sea of Korea* or the *Warsaw Pact*.

4 Discussion of Post Submission Runs

Since the resulting MAPs of our submitted runs, whose parameters had been tuned to fit the GeoCLEF Data of 2005, did not differ substantially, we run additional post experiments trying to isolate the single effects of BRF, NER and weighting and Boolean conjunction. Though we did not submit a German base run without any BRF and weighting, after eliminating an encoding problem such a German base run (no BRF, no NER, similarity based OR) yielded 25.73% MAP using title and description, 23.43% MAP using title, description and narrative and 25.60% MAP when only geographic names from the narratives were used to expand title and description. The respective runs in English resulted in 18.11% MAP vs. 16.27% MAP resp. 20.31% MAP. Table 3 shows MAP-scores of similarity based monolingual experiments systematically varying one factor.

The results show that GeoBRF does lead to considerable better performance for German. Whereas BRF not highlighting (geographic) NEs achieved much lower MAPs than the base runs for both languages, GeoBRF could significantly improve the average precison scores for German. The best BRF parameters were 30 terms from 10 top-ranked documents, which were weighted about half as strong as original query terms. In combination with weighting geographical entities high and other NEs moderately higher than the remaining terms performance further improved by up to 21% over the base run value. The influence of NEs other than geographic in weighting the query before BRF can be neglected (maximum 1% change in MAP) because of the nature of the GeoCLEF topics. The best GeoBRF however did not disregard other NEs, but boosted them about a third as much as geographic ones (kept constant within the runs in table 3). For English even GeoBRF did not or only marginally improve results and only when few documents (5,25) where used. We will examine, if this is related to fewer relevant documents for some topics. Half of the topics had less than 10 relevant documents in the collection and only some had more than 20 relevant documents, which could have reduced the pool of "good" documents to take terms from.

Table 3. Effects of weighting and (Geo)BRF for German and English monolingual runs

Fields	(Geo)NEs	BRF			German MAP	English MAP
		Documents	Terms	GeoBRF		
TD	–	–	–	–	25.73	18.11
TDN	–	–	–	–	23.43	16.27
TD	–	5	25	–	19.72	15.71
TDN	–	5	25	–	22.28	15.47
TD	–	10	20	–	21.82	15.74
TDN	–	10	20	–	21.05	16.84
TD	–	10	30	–	19.25	14.66
TDN	–	10	30	–	21.60	16.41
TD	–	5	25	yes	26.43	19.00
TDN	–	5	25	yes	25.71	18.12
TD	–	10	20	yes	27.36	15.69
TDN	–	10	20	yes	27.04	14.22
TD	–	10	30	yes	30.34	17.44
TDN	–	10	30	yes	28.65	15.53
TD	weighted	–	–	–	27.67	20.38
TDN	weighted	–	–	–	25.65	18.42
TD	weighted	5	25	–	23.59	19.03
TDN	weighted	5	25	–	23.33	18.30
TD	weighted	10	20	–	24.72	18.65
TDN	weighted	10	20	–	24.70	17.70
TD	weighted	10	30	–	24.98	17.88
TDN	weighted	10	30	–	24.03	18.04
TD	weighted	5	25	yes	29.54	18.65
TDN	weighted	5	25	yes	25.65	18.20
TD	weighted	10	20	yes	29.07	14.70
TDN	weighted	10	20	yes	27.90	14.83
TD	weighted	10	30	yes	31.20	17.84
TDN	weighted	10	30	yes	28.72	17.50

For German monolingual retrieval GeoBRF added terms like *Regensburg, Straubing, Wachau* for cities along *the Danube* or *the Rhine* (GC 50), *Lower Saxony, Bremen, Hamburg* for *Northern Germany* (GC 42) or *Mühlheim* and *Gelsenkrichen* for *the Ruhr* (GC 33). A good example for a successful GeoBRF at least regarding the idea of thus obtaining expansion names in English monolingual retrieval was *the Middle East* (GC 37), where GeoNEs like *Syria, Israel, Cairo, Egypt, Lebanon, Gaza, Jerusalem, Beirut* and *Gulf* were added. But, of course, in addition to false classified NEs also other geographic names were added e.g. *Boston* or *Madrid*, which do not belong to the region in question.

We will have to analyse in more detail, for which topics the strategy fits best, i.e. does not expand the query with inadequate geographic names. It seems that using the additional information provided by the narratives aggravates this problem of misleading geographic NEs. Considering the improvement in performance by GeoBRF in contrast to "traditional" BRF it therefore would be worth integrating external resources like a gazetteer to only extract names that are within a certain region (e.g. belong to the same branch within a hierarchy). Methods for finding the geographic scope and focus of documents [cf. 1] could further help selecting the geographically best documents among the top-ranked documents.

But in general the MAPs of the respective runs with or without narratives were close. Since geographic expansion by GeoBRF worked well, this may indicate that the narratives do perhaps not contain the right kind of additional geographic information or at least it is missing for certain topics. To evaluate the expansion through narrative information in combination with Boolean retrieval we ran experiments varying NER, weighting, BRF and GeoBRF always adding geographic NEs to a separate Boolean clause. For German the results were really poor. Best performed a base run with 21.21% MAP (14.62% using narratives), all other runs only achieved about 5 -16% MAP. Traditional BRF with 20 terms from 10 documents did better than any GeoBRF. Therefore only monolingual English runs are considered in table 4.

Table 4. Effects of (Geo)BRF for English monolingual runs using Boolean conjunction

Fields	NEs	BRF			English MAP
		Documents	Terms	GeoBRF	
TD	–	–	–	–	21.23
TDN	–	–	–	–	17.60
TD	weighted	–	–	–	21.23
TDN	weighted	–	–	–	19.41
TD	weighted	10	20	–	18.99
TDN	weighted	10	20	–	20.05
TD	–	5	25	–	18.09
TDN	–	5	25	–	19.49
TD	–	10	30	–	18.95
TDN	–	10	30	–	20.18
TD	–	10	20	yes	18.04
TDN	–	10	20	yes	20.91
TD	–	5	25	yes	14.08
TDN	–	5	25	yes	14.91
TD	–	10	30	yes	13.49
TDN	–	10	30	yes	18.56

Again the best English run did not need NER, weighting or (Geo)BRF and did not use the narrative field. It was even superior to the base run with similarity based searching. The additional weighting of geographic NEs was disastrous to quality, whereas other NEs had no influence at all. Boosting NEs higher did only change the base run, in which the narrative was used.

Regarding BRF the number of documents seems crucial to success; here in contrast to the OR-condition 10 documents were better than 5. With fewer documents and without the narrative especially GeoBRF worsened performance. Thus whereas without BRF or GeoBRF using the narrative information decreased average precisons for English (and German) by up to 7% MAP, BRF and GeoBRF reduced this negative impact substantially and even produced higher MAPs.

Moreover we experimented with first running a ranking query and then adding geographic entities from the BRF as a Boolean clause. For both German and English this generally brought worse results than the base runs. Though it could improve MAP by 2% (20.20% MAP) for English and MAP by 3% (28.43% MAP) for German by weighting NEs in the original OR-query and then expanding via GeoBRF with the GeoNEs combined via AND. The optimal number of (geographic) terms to add in the BRF was thereby 25 from the 5 top-ranked documents for English and 30 terms from 15 documents weighted again about half as much as terms from the original query. Here expanding with narrative information led to poor results.

Summing up, since we could not find a substantial positive impact of additional geographic information, the question of geographic expansion and Boolean retrieval remains to be examined. Particularly different kinds of geographic information to use for expansion must be tested and heuristics must then be established. First experiments with manual expansion of German queries combined with the Boolean approach support the assumption that the narrative information may not be adequate for the German collection. Using e.g. names of cities for regions like the *Ruhrgebiet* and for GC 26; names of countries and their capitals resp. their five biggest cities, MAPs were about 4 to 10% (best: 24.81% MAP) better than runs with using the narratives for expansion.

5 Conclusion and Outlook

Query expansion by adding particularly geographic NEs from the top-ranked documents within the process of BRF has shown to substantially improve retrieval performance, especially in combination with weighting geographic NEs in the original query very high. With respect to a Boolean approach to GIR our results are still mixed, probably also due to insufficient NER. Therefore we are currently working on improving the NER by training advanced models and testing the fusion of machine learning with other approaches like gazetteer-look up and handcrafted rules. The integration of a gazetteer for automatic query expansion is planned, but priority is given to explore different heuristics for such an expansion, also to further clarify the adequacy of Boolean retrieval. Factors like proximity, connection (economic, traditional, via language) and especially for the domain of newspaper articles also certain criteria of news selection like elite nation, sensationalism etc. may influence the publicity of location names and thus the terms resp. the granularity used for their reference.

To pursue the idea of geographic BRF techniques for disambiguation and the identification of the geographical scope of a document like in [1] should be integrated to only add geographically related information. Such approaches would also be crucial for using Wikipedia as an additional external resource, by which geographic references not captured in gazetteers should be expandable. Moreover we will also enhance the system by enabling expansion with synonyms e.g. via the linguistic internet service *Wortschatz Leipzig*[7] for German. A detailed analysis of system performance on the individual topics however has to show the feasibility of a text based approach without any Geo-Coding and –matching. Some topics of this year would seem to demand elaborated spatial methods or techniques of natural language processing (NLP) to identify place types e.g. cities and negations (e.g. *except the republics of the former USSR*). However, it should be noted that some of these topics might not mirror realistic user needs.

References

[1] Amitay, Einat; Har'El, Nadav; Sivan, Ron; Soffer, Aya (2004): *Web-a-Where: Geotagging Web content*. In: SIGIR 2004: Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. July 2004, Sheffield, UK. ACM, pp. 273-280.

[2] Chaves, Marcirio Silveira; Martins, Bruno; Silva, Mário J. (2005): *Challenges and Resources for Evaluating Geographical IR*. In: Proceedings of the 2nd International Workshop on Geographic Information Retrieval, CKIM 2005. Nov. 2005, Bremen, Germany. pp. 65-69.

[3] Clough, Paul (2005): *Extracting Metadata for Spatially-Aware Information Retrieval on the Internet*. In: Proceedings of the 2nd International Workshop on Geographic Information Retrieval, CKIM 2005. Nov. 2005, Bremen, Germany. pp 25-30.

[4] Clough, Paul; Joho, Hideo; Purves, Ross (2005): *Identifying imprecise regions for geographic information retrieval using the web*. In: Proceedings of the GIS RESEARCH UK 13th Annual Conference. Glasgow, UK, 2005. pp. 313-318.

[5] Ageno, Alicia; Ferrés, Daniel; Rodríguez, Horacio (2005): *The GeoTALP-IR System at GeoCLEF-2005: Experiments Using a QA-based IR System, Linguistic Analysis, and a Geographical Thesaurus*. In: Working Notes of the 6th Workshop of the Cross-Language Evaluation Forum, CLEF. Sep. 2005, Vienna, Austria. http://www.clef-campaign.org/2005/ working_notes/.

[6] Gey, Fredric; Larson, Ray; Sanderson, Mark; Joho, Hideo; Clough, Paul (2005): *GeoCLEF: the CLEF 2005 Cross-Language Geographic Information Retrieval Track*. In: Working Notes of the 6th Workshop of the Cross-Language Evaluation Forum, CLEF. Sep. 2005, Vienna, Austria. http://www.clef-campaign.org/ working_notes/.

[7] Gey, Frederic; Petras, Vivien (2005): *Berkeley2 at GeoCLEF: Cross-Language Geographic Information Retrieval of German and English Documents*. In: Working Notes of the 6th Workshop of the Cross-Language Evaluation Forum, CLEF. Sep. 2005, Vienna, Austria. http://www.clef-campaign.org/ working_notes/.

[8] Hackl, René; Mandl, Thomas; Womser-Hacker, Christa (2005): *Mono- and Cross-lingual Retrieval Experiments at the University of Hildesheim*. In: Peters, Carol; Clough, Paul; Gonzalo, Julio; Kluck, Michael; Jones, Gareth; Magnini, Bernard (eds): Multilingual Information Access for Text, Speech and Images: Results of the Fifth CLEF Evaluation Campaign. Berlin et al.: Springer [LNCS 3491] pp. 165-169.

[9] Larson, Ray R. (2005): *Chesire II at GeoCLEF: Fusion and Query Expansion for GIR*. In: Working Notes of the 6th Workshop of the Cross-Language Evaluation Forum, CLEF. Sep. 2005, Vienna, Austria. http://www.clef-campaign.org/2005/working_notes/

[10] Mandl, Thomas; Schneider, René; Strötgen, Robert (2006): *A Fast Forward Approach to Cross-lingual Question Answering for English and German*. In: Gey, Fredric C.; Gonzalo, Julio; Jones, Gareth J.F; Kluck, Michael; Magnini, Bernardo; Müller, Henning; Peters, Carol; de Rijke, Maarten (eds.). Accessing Multilingual Information Repositories: 6th Workshop of the Cross-Language Evaluation Forum, CLEF. Sep. 2005, Vienna, Austria. Revised Selected Papers. LNCS 4022; Springer, pp. 332-336.

[7] http://wortschatz.uni-leipzig.de/