Applying NLP Techniques for Query Reformulation to Information Retrieval with Geographical References

José M. Perea-Ortega[1], Miguel A. García-Cumbreras[2], and L. Alfonso Ureña-López[2]

[1] Languages and Information Systems Department, University of Sevilla
E.T.S. Ingeniería Informática, Avda. Reina Mercedes s/n, 41012, Sevilla, Spain
jmperea@us.es
[2] Computer Science Department, University of Jaén
Escuela Politécnica Superior, Campus Las Lagunillas s/n, 23071, Jaén, Spain
{magc,laurena}@ujaen.es

Abstract. Geographic Information Retrieval (GIR) is an active and growing research area that focuses on the retrieval of textual documents according to a geographical criteria of relevance. However, since a GIR system can be treated as a traditional Information Retrieval (IR) system, it is important to pay attention to finding effective methods for query reformulation. In this way, the search results will improve their quality and recall. In this paper, we propose different Natural Language Processing (NLP) techniques of query reformulation related to the modification and/or expansion of both parts thematic and geospatial that are usually recognized in a geographical query. We have evaluated each of the reformulations proposed using GeoCLEF as an evaluation framework for GIR systems. The results obtained show that all proposed query reformulations retrieved relevant documents that were not retrieved using the original query.

Keywords: Geographic query reformulation, Geographic Information Retrieval, Query expansion, GeoCLEF.

1 Introduction

In the Information Retrieval (IR) field [2], the approach based on the modification of the user query to improve the quality of the IR results is known as query reformulation. The aim of such process is to satisfy the user information need, usually improving the quality and recall of the results obtained using the original user query. This feature is explicitly supported by some search engines suggesting related queries or providing different completions of the initial user query. Moreover, other search engines also support query reformulation in an implicit manner, by expanding the original query with terms related to their keywords, for example.

T. Washio and J. Luo (Eds.): PAKDD 2012 Workshops, LNAI 7769, pp. 57–69, 2013.

Geographic Information Retrieval (GIR) is an active and growing research area that focuses on the retrieval of textual documents according to a geographical criteria of relevance. For this reason, GIR is considered as an extension of the field of IR. Specifically, GIR is concerned with improving the quality of geographically-specific information retrieval, focusing on access to unstructured documents [10,13]. The IR community has primarily been responsible for research in the GIR field, rather than the Geographic Information Systems (GIS) community. The type of query in a IR engine is based usually on natural language, in contrast to the more formal approach common in GIS, where specific geo-referenced objects are retrieved from a structured database. In a GIR system, a geographic query can be structured as a triplet of <*theme*><*spatial relationship*><*location*>, where <*theme*> is the main subject of the query, <*location*> represents the geographical scope of the query and <*spatial relationship*> determines the relationship between the subject and the geographical scope. For example, the triplet for the geographical query "*airplane crashes close to Russian cities*" would be <*airplane crashes*><*close to*><*Russian cities*>. Thus, a search for "*castles in Spain*" should return not only documents that contain the word "*castle*", but also those documents which have some geographical entity related to Spain.

Since a GIR system can be treated as a traditional search engine (the results for a query are displayed as a ranked list), it is important to pay attention to finding effective methods for query reformulation. These methods can take into account both lexical-syntactic features and geographical aspects. In this way, the search results will improve their quality and recall. The objective of this paper is to evaluate several geographic query reformulations for the GIR task, considering that a GIR system can perform as a IR system. To carry out this evaluation, we have used the most important evaluation framework in this context: GeoCLEF[1] [7,14].

The remainder of this paper is structured as follows: in Section 2, the most important works related to the geographic query reformulation in GIR are expounded; in Section 3, we describe the GIR system used for the experiments; Section 4 presents the main features of the query reformulations proposed; in Section 5, we describe briefly the evaluation framework; in Section 6 and Section 7, the experiments carried out and an analysis of the results are presented; finally, in Section 8, we draw some conclusions and future work is expounded.

2 Related Work

Jansen et al. [9] define the concept of query reformulation as the process of altering a given query in order to improve search or retrieval performance. Sometimes, query reformulation is applied automatically by search engines as with *relevance feedback* technique. It is a method that allows users to judge whether a document is relevant or not, so that automatic rewritings can be generated depending on it. At other times, query reformulation is carried out analysing

[1] `http://ir.shef.ac.uk/geoclef/`

the top retrieved documents without the user's intervention, taking into account term statistics. However, it has been found that users rarely utilize the relevance feedback options [19] and usually reformulate their needs manually [1].

The focus of this paper is geographic queries. According to Gravano [8], search engines are criticised because of their ignorance to the geographical constraints on users' queries and, therefore, retrieve less relevant results. This could be attributed to the way search engines handle queries in general as they adopt a keywords matching approach without spatially inferring the scope of the geographic terms. However, it shall be noted that a number of services to deal with this issue have recently been proposed in major search engines, but not in the general purpose tools.

Several authors have studied what users are looking for when submitting geographic queries [18,6,11]. One of the main conclusions of these studies is that the structure of geographic queries consists of thematic and geographical parts, with the geo-part occasionally containing spatial or directional terms. From a geographical point of view, Kohler [12] provides a research about geo-reformulation of queries. She concludes that the addition of more geo terms in the query is commonly used to differentiate between places that share the same name. This is also known as query expansion using geographic entities.

In the literature, we can find various works that have addressed the spatial query expansion. Cardoso et al. [4] present an approach for geographical query expansion based on the use of feature types, readjusting the expansion strategy according to the semantics of the query. Fu et al. [5] propose an ontology-based spatial query expansion method that supports retrieval of documents that are considered to be spatially relevant. They improve search results when a query involves a fuzzy spatial relationship, showing that proposed method works efficiently using realistic ontologies in a distributed spatial search environment. Buscaldi et al. [3] use WordNet[2] during the indexing phase by adding the synonyms and the holonyms of the encountered geographical entities to each documents index terms, proving that such method is effective. Finally, Stokes et al. [20] conclude that significant gains in GIR will only be made if all query concepts (not just geospatial ones) are expanded.

Our work could be positioned within those works that treat geographical part as textual terms, i.e., from a Natural Language Processing (NLP) point of view, exclusively. For this reason, the proposed query reformulations are based on expansions of the thematic and geographical parts detected in a geographical query, using synonyms and geospatial terms related with the keywords and geographical entities found in the query.

3 The SINAI-GIR Architecture

In this Section we describe an example of a GIR system. Specifically, we have used our own GIR system called SINAI-GIR [17]. GIR systems are usually composed of three main stages: preprocessing of the document collection and queries,

[2] `http://wordnet.princeton.edu/`

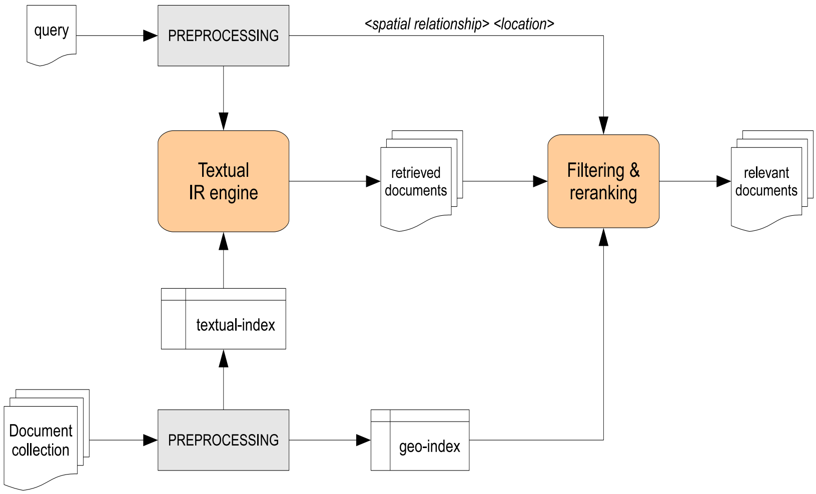

Fig. 1. Overview of the SINAI-GIR system

textual-geographical indexing and searching and, finally, reranking of the retrieved results using a particular relevance formula that combines textual and geographical similarity between the query and the document retrieved. This general architecture is shown in Figure 1.

With respect to the document collection processing, it is based on detecting all the geographical entities in each document and generating a geo-index with them. In this phase, the stop words are removed and the stem of each word is taken into account. We have used our own Named Entity Recognition (NER) tool to detect geographical entities. It is called GeoNER [16] and it is based on external knowledge resources such as GeoNames[3] and Wikipedia.

Regarding query processing, each query is preprocessed and analyzed, identifying the geographical scope and the spatial relationship that may contain. It also involves specifying the triplet explained in Section 1, which will be used later during the filtering and reranking process. To detect such triplet, we have used a Part Of Speech tagger (POS tagger) like TreeTagger[4], taking into account some lexical syntactic rules such as *preposition + proper noun*, for example. Moreover, the stop words are removed and the Snowball stemmer[5] is applied to each word of the query, except for the geographical entities. During the text retrieval

[3] GeoNames is a geographical database covers all countries and contains over eight million placenames that are available for download free of charge. `http://www.geonames.org`

[4] TreeTagger v.3.2 for Linux. Available in `http://www.ims.uni-stuttgart.de/projekte/corplex/TreeTagger/DecisionTreeTagger.html`

[5] Available in `http://snowball.tartarus.org`

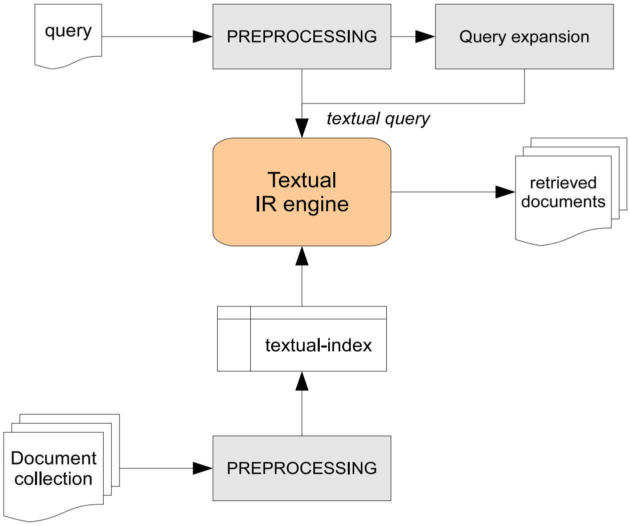

Fig. 2. Architecture of the GIR system employed to carry out the experiments

process, we obtain 1,000 documents for each query. We have used Terrier[6] as a search engine. According to a previous work [15], it was shown that Terrier is one of the most used IR tools in IR systems in general and GIR systems in particular, obtaining promising results. The weighting scheme used has been *inL2*, which is implemented by default in Terrier. This scheme is the Inverse Document Frequency (IDF) model with Laplace after-effect and normalization two. As a final step, each preprocessed query (including their geographical entities) is run against the search engine. The retrieved documents are filtered and reranked, setting in the last positions those documents that do not match with the geographical scope detected in the query. By contrast, those documents that fit the geographical scope detected, are set in the first positions.

Although GIR systems usually apply a geo-reranking process after the IR module (as can be seen in Figure 1), it is important to note that this process is not necessary for this work particularly, because we are interested in analyzing the behaviour of each proposed query reformulation from an IR point of view, i.e. evaluating their precision and recall scores without any reranking process that applies a geographic reasoning. Therefore, the architecture employed to carry out the experiments of this work follows the similar approach that is applied in traditional information retrieval systems but considering query reformulations, as shown in Figure 2.

[6] Version 2.2.1, available in `http://terrier.org`

4 Query Reformulations Proposed

Several query reformulations for the GIR task are analyzed in this work. They try to use both the thematic part and the geographical scope detected in the query. The objective of these query reformulations is to improve the retrieval process trying to find relevant documents that are not retrieved using the original query. Starting from the preprocessed original query, we have generated the following query reformulations:

- QR1: the geographical scope is removed, leaving only the thematic part of the original query.
- QR2: the thematic part is expanded, repeating its terms. In this way, we try to give more importance to the thematic part than the geo-part.
- QR3: the thematic part is expanded using only synonyms of the keywords detected in the thematic part of the query. We have considered as keywords the nouns recognized in such part. WordNet was used as external resource in order to extract the synonyms for each keyword.
- QR4: the geographical part is expanded using only synonyms of the geographical scope detected in the query. These synonyms were extracted from the GeoNames database.
- QR5: the geographical part is expanded using locations or places that match with the geographical scope and the spatial relationship detected in the query. Like the previous query reformulation, GeoNames was used as geographical knowledge base.
- QR6: the thematic and geographical parts are expanded, combining the QR3 and QR5 reformulations.

Table 1. Example of query reformulations generated for the query "*Visits of the American president to Germany*"

Reformulation	Text of the query
original	`visit American presid Germany`
QR1	`visit American presid`
QR2	`visit American presid visit American presid Germany`
QR3	`#and(#or(visit meet stay ) American presid Germany)`
QR4	`#and(visit American presid #or(Germany #3(Federal Republic of Germany) Deutschland FRG ) )`
QR5	`#and(visit American presid #or(Germany Berlin Hamburg Muenchen Koeln #2(Frankfurt am Main) Essen ) )`
QR6	`#and(#or(visit meet stay ) of the American presid ) #or(Germany Berlin Hamburg Muenchen Koeln #2(Frankfurt am Main) Essen )`

Table 1 shows an example of the different query reformulations generated for the query "*Visits of the American president to Germany*". As can be seen, QR2 and QR3 are query reformulations that expand only the thematic part of the queries and, on the other hand, QR4 and QR5 expand only the geographical part of them. Finally, QR6 can be considered a combination of expansions using both parts.

5 GeoCLEF: The Evaluation Framework

In order to evaluate the proposed query reformulations, we have used the GeoCLEF framework [7,14], an evaluation forum for GIR systems held between 2005 and 2008 under the CLEF[7] conferences. GeoCLEF provides a document collection that consists of 169,477 documents, composed of stories and newswires from the British newspaper *Glasgow Herald* (1995) and the American newspaper *Los Angeles Times* (1994), representing a wide variety of geographical regions and places. On the other hand, there are a total of 100 textual queries or topics provided by GeoCLEF organizers (25 per year). They are composed of three main fields: *title* (T), *description* (D) and *narrative* (N). For the experiments carried out in this work, we have only taken into account the *title* field. Some examples of GeoCLEF topics are: "*vegetable exporters of Europe*", "*forest fires in north of Portugal*", "*airplane crashes close to Russian cities*" or "*natural disasters in the Western USA*".

Regarding the evaluation measures used, results are evaluated using the relevance judgements provided by the GeoCLEF organizers and the TREC evaluation method. The evaluation has been accomplished by using the Mean Average Precision (MAP), Recall (R) and Precision at n (P@n). The MAP measure computes the average precision over all queries. The average precision is defined as the mean of the precision scores obtained after each relevant document is retrieved, using zero as the precision for relevant documents that are not retrieved. Recall is a measure of the extent to which relevant documents are found or retrieved. Recall is 1.0 when every relevant document is retrieved. Finally, Precision at n is the precision at the number of n relevant documents in the collection for the query. Precision is the fraction of the relevant documents divided by the total number of documents retrieved. Therefore, if P@n is 1.0, it means a perfect relevance ranking and a perfect recall at n documents retrieved.

6 Experiments and Results

The different results obtained using each query reformulation (QR) along with the result obtained using the original query are shown in Table 2. In such table, we show the average score of precision at the 5, 10 and 100 first documents retrieved, recall (R) and MAP for each query reformulation proposed. Although none of the proposed QRs improve the MAP score obtained using the original query, it is interesting to note that QR2 (the thematic part is expanded, repeating its terms) achieves the best P@10 score in three of the four topic sets.

[7] `http://www.clef-initiative.eu/`

Table 2. Evaluation results obtained for each query reformulation proposed

Topic Set	QR	P@5	P@10	P@100	R	MAP
2005	original	**0.5520**	0.4560	**0.1904**	**0.8364**	**0.3514**
	QR1	0.2640	0.2560	0.1260	0.6748	0.1638
	QR2	0.5200	**0.4920**	0.1840	0.8276	0.3353
	QR3	0.3680	0.3160	0.1400	0.7596	0.2035
	QR4	0.3120	0.2800	0.1212	0.6552	0.2242
	QR5	0.1440	0.1240	0.0772	0.5624	0.0952
	QR6	0.1600	0.1480	0.0780	0.5692	0.0942
2006	original	**0.2400**	0.1920	**0.0716**	**0.7288**	**0.2396**
	QR1	0.0560	0.0640	0.0252	0.4604	0.0615
	QR2	0.2320	**0.2040**	0.0664	0.6796	0.2314
	QR3	0.1440	0.1400	0.0604	0.7356	0.1419
	QR4	0.1920	0.1720	0.0636	0.6984	0.2064
	QR5	0.2240	0.1840	0.0612	0.6524	0.1811
	QR6	0.1840	0.1760	0.0580	0.6772	0.1486
2007	original	**0.3040**	**0.2560**	**0.1188**	**0.7156**	**0.2311**
	QR1	0.1600	0.1320	0.0796	0.4452	0.1255
	QR2	0.2640	0.2120	0.1072	0.6656	0.1871
	QR3	0.2000	0.1800	0.0884	0.6284	0.1774
	QR4	0.2160	0.2000	0.1020	0.6608	0.1687
	QR5	0.2240	0.2000	0.0928	0.6720	0.1874
	QR6	0.2240	0.2040	0.0836	0.6344	0.1763
2008	original	**0.3760**	**0.2680**	0.1104	**0.7368**	**0.2484**
	QR1	0.1760	0.1400	0.0928	0.5996	0.1301
	QR2	0.3440	**0.2680**	**0.1124**	0.7196	0.2381
	QR3	0.2960	0.2320	0.1024	0.6884	0.1972
	QR4	0.2640	0.1960	0.0924	0.6404	0.1619
	QR5	0.2720	0.2040	0.0964	0.6984	0.1906
	QR6	0.2720	0.2280	0.0948	0.7028	0.2028

At this point, we wonder if the QRs proposed were really retrieving relevant documents that the original query was not retrieving. Using the relevance judgements provided by the GeoCLEF organizers, we get the relevant documents retrieved by each QR that were not retrieved by the original query, as shown in Figure 3 and Table 3. The total number of documents retrieved was always 1,000. While in Figure 3 we can compare the behaviour of each query reformulation for the different topic sets regarding the number of relevant documents that were not retrieved by the original query, Table 3 presents these results numerically making a comparison with those obtained using the original query. It is also shown the total number of relevant documents for each topic set.

Analyzing these results in general, we can observe that all proposed query reformulations always retrieved relevant documents that were not retrieved using the original query. This does not mean that the proposed query reformulations achieve higher MAP scores than those obtained by the original query (see Table 2). The main reason for this behaviour is focused on the ranking process. In this

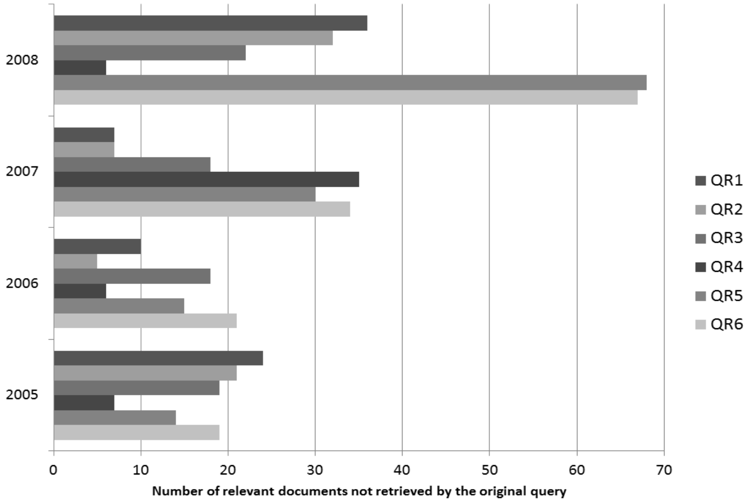

Fig. 3. Comparison of the number of relevant documents retrieved by each query reformulation that were not retrieved by the original query

Table 3. Number of relevant documents retrieved by each query reformulation compared with those obtained using the original query

Query set	Total num relevant docs	Num of relevant doc retrieved						
		original	QR1	QR2	QR3	QR4	QR5	QR6
2005	1028	**908**	735	895	813	706	579	583
2006	378	**284**	160	251	264	280	272	255
2007	650	**543**	391	521	489	483	493	464
2008	747	588	529	**597**	577	480	539	542

experiments we have not applied any spatial ranking process, only has been used the ranking provided by the search engine that does not employ any geographic reasoning. Another reason is that none of the query reformulations outperform the total number of relevant documents retrieved by the original query (except QR2 for the 2008 topic set), as can be seen in Table 3.

7 Analysis and Discussion

Following a general analysis, it is interesting to note the behaviour of the reformulations related to the geographical expansion (QR4 and QR5). Specifically, QR5 achieves a remarkable difference using the 2008 topic set with a total of 68 relevant documents not retrieved by the original query. In fact, this means that of 159 relevant documents not retrieved by the original query using the 2008

Table 4. MAP average results according to the query type

query type	MAP						
	original	**QR1**	**QR2**	**QR3**	**QR4**	**QR5**	**QR6**
Part-to relationship	**0.2582**	0.0819	0.2323	0.1686	0.1781	0.1411	0.1498
Adjacent-to relationship	**0.2130**	0.0861	0.2061	0.1273	0.1179	0.1126	0.0881
Non-geographic	**0.3771**	0.3510	0.3742	0.2974	0.3342	0.3342	0.2607

Table 5. P@10 average results according to the query type

query type	P@10						
	original	**QR1**	**QR2**	**QR3**	**QR4**	**QR5**	**QR6**
Part-to relationship	**0.2986**	0.1114	0.2971	0.2200	0.2114	0.1600	0.1829
Adjacent-to relationship	0.1813	0.1000	**0.1875**	0.1125	0.0938	0.1063	0.1375
Non-geographic	0.3929	0.3857	**0.4000**	0.3214	0.3500	0.3500	0.2786

topic set (747-588), 42.77% of them are retrieved using the QR5 reformulation. Another example occurs with QR4 that obtains the highest value for the 2007 topic set, retrieving 32.71% of the relevant documents not retrieved using the original query. On the other hand, the reformulations related to the thematic expansion (QR2 and QR3) also achieve good results in general, as can be seen for the 2005 and 2006 topic sets. All this makes that the reformulation that combines the QR3 and QR5 reformulations (QR6) also obtain good results, as shown for all topic sets. Finally, QR1 achieves the best score for the 2005 topic set, so the idea of removing the geographical part in the original query can sometimes be a good strategy. This may sound a little strange when we are working on GIR, but we have to take into account that sometimes a query can be considered as a geographic query because it contains a geographical term, but really it is not. For instance, the query "*Japanese rice imports*" might seem a geographic query because it contains the term "*Japanese*", but really it does not impose any geographical constraint.

In order to carry out a more in-depth analysis regarding the distinctive features each QR has, we will use the classification type given by Cardoso and Silva regarding spatial relationships [4]. They distinguish two main types (*part-of* and *adjacent-to*) in order to drive the query expansion strategies according to the proper relationships contained in the geographical ontology used for that purpose. *Part-of* relationships (for example *in*, *of*, *on the*, *at*, etc.) are the most common spatial relationships found on geographical queries [12], denoting that the user is interested on documents inside the boundaries of the given scope of interest. Adjacent-to relationships denote proximity (for example *around*, *next to*, *within X km of*, etc.) and their semantic may have distinct interpretations [5]. According to this classification and taking into account that a geographic query can be considered as a non-geographic query despite contain a geographic entity, we classified manually the 100 queries provided by GeoCLEF, resulting

Table 6. Number of relevant documents retrieved according to the query type

query type	Number of relevant documents retrieved						
	original	QR1	QR2	QR3	QR4	QR5	QR6
Part-to relationship	**1643**	1234	1623	1507	1336	1227	1231
Adjacent-to relationship	**284**	206	248	244	220	263	224
Non-geographic	**396**	375	393	392	393	393	389

14 as non-geographic, 16 as *adjacent-to* type and the remaining (70) as *part-of* type. Therefore, most of the GeoCLEF queries (70%) are considered as *part-of* type.

According to that classification and based on the results shown in Table 4, Table 5 and Table 6, we can observe some findings. As expected, the behaviour of the QR1 is good in general for those queries considered as non-geographic, although that reformulation type does not improve any of the results obtained for the original query on average. In view of the obtained results on average, we can not draw a clear conclusion about when is more desirable to apply one reformulation type according to the type of the spatial relationship detected in the query. However it is interesting to note the good performance of the QR2 for the P@10 measure in general, and for the *adjacent-to* and *non-geographic* query types in particular. This means that repeat the keywords of the thematic part detected in a query could be a good strategy in order to obtain more relevant documents in these systems, particularly when we submit *non-geographic* or *adjacent-to* query types. This behaviour can be explained because by repeating the keywords in the thematic part we are reinforcing the importance of the information need provided by the user in the query when the geographical constraint is not so important.

8 Conclusions and Further Work

In this paper we propose different NLP techniques of query reformulation related to the modification and/or expansion of both parts thematic and geospatial that are usually recognized in a geographical query. We have evaluated each of the reformulations proposed using GeoCLEF as an evaluation framework for GIR systems. This evaluation has been carried out from an IR point of view, that is, without taking into account any geo-reranking procedure after the retrieval process. The results obtained show that all proposed query reformulations retrieved relevant documents that were not retrieved using the original query although these did not improve the results obtained using the original query on average. We carried out a brief analysis according to the two main types of spatial relationships that can be recognized in a geographical query, but it did not provide us a clear conclusion. However, we noted that repeat the keywords of the thematic part detected in a query could be a good strategy in order to obtain more relevant documents in these systems, particularly when we submit *non-geographic* or *adjacent-to* query types.

For future work, we will study in depth when is more suitable to apply these techniques in a GIR system depending on the type of the query and providing a fusion method for collecting those relevant documents retrieved by each query reformulation that were not retrieved using the original query. Then we will work on the spatial reranking process after this fusion in order to sort the final list of documents according to the two criteria of relevance in these systems: thematic and geographical.

Acknowledgments. This work has been partially funded by the European Commission under the Seventh (FP7-2007-2013) Framework Programme for Research and Technological Development through the FIRST project (FP7-287607). This publication reflects the views only of the author, and the Commission cannot be held responsible for any use which may be made of the information contained therein. It has been also partially supported by a grant from the Fondo Europeo de Desarrollo Regional (FEDER) through the TEXT-COOL 2.0 project (TIN2009-13391-C04-02) from the Spanish Government.

References

1. Anick, P.: Using terminological feedback for web search refinement: a log-based study. In: SIGIR 2003: Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Informaion Retrieval, pp. 88–95. ACM, New York (2003)
2. Baeza-Yates, R.A., Ribeiro-Neto, B.: Modern Information Retrieval. Addison-Wesley Longman Publishing Co., Inc., Boston (1999)
3. Buscaldi, D., Rosso, P., Arnal, E.S.: Using the WordNet Ontology in the GeoCLEF Geographical Information Retrieval Task. In: Peters, C., Gey, F.C., Gonzalo, J., Müller, H., Jones, G.J.F., Kluck, M., Magnini, B., de Rijke, M. (eds.) CLEF 2005. LNCS, vol. 4022, pp. 939–946. Springer, Heidelberg (2006)
4. Cardoso, N.: Query expansion through geographical feature types. In: Purves, R., Jones, C. (eds.) GIR, pp. 55–60. ACM (2007)
5. Fu, G., Jones, C.B., Abdelmoty, A.I.: Ontology-Based Spatial Query Expansion in Information Retrieval. In: Meersman, R., Tari, Z. (eds.) OTM 2005. LNCS, vol. 3761, pp. 1466–1482. Springer, Heidelberg (2005)
6. Gan, Q., Attenberg, J., Markowetz, A., Suel, T.: Analysis of geographic queries in a search engine log. In: Proceedings of the First International Workshop on Location and the Web, pp. 49–56. ACM, Beijing (2008)
7. Gey, F.C., Larson, R.R., Sanderson, M., Joho, H., Clough, P., Petras, V.: GeoCLEF: The CLEF 2005 Cross-Language Geographic Information Retrieval Track Overview. In: Peters, C., Gey, F.C., Gonzalo, J., Müller, H., Jones, G.J.F., Kluck, M., Magnini, B., de Rijke, M., Giampiccolo, D. (eds.) CLEF 2005. LNCS, vol. 4022, pp. 908–919. Springer, Heidelberg (2006)
8. Gravano, L., Hatzivassiloglou, V., Lichtenstein, R.: Categorizing web queries according to geographical locality. In: Proceedings of the 12th International Conference on Information and Knowledge Management, pp. 325–333 (2003)
9. Jansen, B.J., Booth, D.L., Spink, A.: Patterns of query reformulation during web searching. JASIST 60(7), 1358–1371 (2009)

10. Jones, C.B., Purves, R.S.: Geographical information retrieval. International Journal of Geographical Information Science 22(3), 219–228 (2008)
11. Jones, R., Zhang, W.V., Rey, B., Jhala, P., Stipp, E.: Geographic intention and modification in web search. International Journal of Geographical Information Science 22(3), 229–246 (2008)
12. Kohler, J.: Analysing search engine queries for the use of geographic terms. Master's thesis, University of Sheffield - United Kingdom (2003)
13. Larson, R.: Geographic information retrieval and spatial browsing. In: Smith, Gluck, M. (eds.) Geographic Information Systems and Libraries: Patronsand Mapsand and Spatial Information, pp. 81–124 (1996)
14. Mandl, T., Carvalho, P., Di Nunzio, G.M., Gey, F., Larson, R.R., Santos, D., Womser-Hacker, C.: GeoCLEF 2008: The CLEF 2008 Cross-Language Geographic Information Retrieval Track Overview. In: Peters, C., Deselaers, T., Ferro, N., Gonzalo, J., Jones, G.J.F., Kurimo, M., Mandl, T., Peñas, A., Petras, V. (eds.) CLEF 2008. LNCS, vol. 5706, pp. 808–821. Springer, Heidelberg (2009)
15. Perea-Ortega, J.M., García-Cumbreras, M.Á., García-Vega, M., Ureña-López, L.A.: Comparing Several Textual Information Retrieval Systems for the Geographical Information Retrieval Task. In: Kapetanios, E., Sugumaran, V., Spiliopoulou, M. (eds.) NLDB 2008. LNCS, vol. 5039, pp. 142–147. Springer, Heidelberg (2008)
16. Perea-Ortega, J.M., Martínez-Santiago, F., Montejo-Ráez, A., Ureña-López, L.A.: Geo-NER: un reconocedor de entidades geográficas para inglés basado en GeoNames y Wikipedia. Sociedad Española para el Procesamiento del Lenguaje Natural (SEPLN) 43, 33–40 (2009)
17. Perea-Ortega, J.M., Ureña-López, L.A., García-Vega, M., García-Cumbreras, M.A.: Using Query Reformulation and Keywords in the Geographic Information Retrieval Task. In: Peters, C., Deselaers, T., Ferro, N., Gonzalo, J., Jones, G.J.F., Kurimo, M., Mandl, T., Peñas, A., Petras, V. (eds.) CLEF 2008. LNCS, vol. 5706, pp. 855–862. Springer, Heidelberg (2009)
18. Sanderson, M., Kohler, J.: Analyzing geographic queries. In: Proceedings Workshop on Geographical Information Retrieval SIGIR (2004)
19. Spink, A., Jansen, B.J., Ozmultu, C.H.: Use of query reformulation and relevance feedback by excite users. Internet Research: Electronic Networking Applications and Policy 10(4), 317–328 (2000)
20. Stokes, N., Li, Y., Moffat, A., Rong, J.: An empirical study of the effects of nlp components on geographic ir performance. International Journal of Geographical Information Science 22(3), 247–264 (2008)