

The University of Lisbon at GeoCLEF 2007

Nuno Cardoso, David Cruz, Marcirio Chaves and Mário J. Silva
University of Lisbon, Faculty of Sciences
1749-016 Lisboa, Portugal
{ncardoso, dcruz, mchaves, mjs} @ xldb.di.fc.ul.pt

Abstract

This paper reports the participation of the XLDB Group from the University of Lisbon at the 2007 GeoCLEF task. We adopted a novel approach for GIR, focused on handling geographic features and feature types on both queries and documents, generating geographic signatures with multiple geographic concepts as a scope of interest. We experimented new query expansion and text mining strategies, relevance feedback approaches and geographic score metrics. In the paper we introduce the new approach, discuss the experiments and analyse the obtained results.

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

Evaluation, Geographic IR, Text Mining, Geographic Relevance, GeoCLEF

1 Introduction

This paper presents the participation of the XLDB Group from the University of Lisbon at the 2007 GeoCLEF task. We experimented novel strategies for geographic query expansion, text mining, relevance feedback and geographic score metrics in a renewed GIR system. The motivation for this work derived from the results obtained in last year's participation, which revealed limitations on our previous GIR model [9]:

- We focused on capturing and handling placenames and associated features from queries and documents for our geographic reasoning, and ignored important geographic information, such as spatial relationships and *feature types*. Feature types, such as *cities*, *mountains* or *airports*, play an important role on the definition of the geographic relevance criteria of queries. GeoCLEF topics also convey this idea: 13 out of the 25 topics of the Portuguese subtask of the 2007 edition of GeoCLEF contained feature types on the topic's title.
- Typical GeoIR systems rely on text mining methods to capture and disambiguate *geonames* present in the text, so that geographic scopes can be inferred for each document. These methods typically involve geoname grounding into *geographic concepts* included in a *geographic ontology*, and disambiguation of hard cases through reasoning based on other geonames extracted from the text [14]. We used this text mining approach in our past GeoCLEF participations [2, 9]. The mining process was

finalized by a graph-ranking algorithm, that analysed the captured features and assigned one single encompassing scope per document [10]. This strategy is derived from the «one scope per discourse» assumption [7], spanned to a full document. The assumption of taking the unit of discourse to the document level revealed to be too restrictive in some cases, and highly vulnerable to incorrectly assigned scopes. We observe that generic scopes were being assigned to documents with geonames that do not correspond to adjacent areas. For example, a document describing a football match between Portugal and Hungary, may have the common ancestor node (Europe) as a very strong candidate final scope.

This year, we decided to challenge some of the underlying assumptions of the GIR model used in the previous year, and tested a new approach. We introduced significant changes in the assembled GIR system, both on the query and on the document sides, to see if they could effectively tackle the limitations detected on the past GIR system. The improvements have been introduced at three levels:

Query Processing: We have rebuilt the query processing modules so that all geographic information present on a query is captured and subject to proper geographic query expansion. We gave special attention to feature types and spatial relationships, as guides for the geographic query expansion [3].

Text Mining: We decided to narrow the discourse context to the sentence level. We now generate what we call a *geographic signature* for each document, which is a list of geographic concepts that characterize a document, allowing each document to have several geographic contexts.

Geographic Ranking: As the new text mining approach generates a geographic signature for each document (D_{Sig}), and the geographic query expansion module generates a geographic signature for the query (Q_{Sig}), the geographic ranking step now has the burden of evaluating relevance considering queries and documents that contain multiple geographic concepts as a scope. In 2006, our similarity metric compared the (single) scope of a document against the (single) scope of a query. This year, we had to handle each of the features in the geographic signatures as part of a scope and compute a metric accounting for all concepts in the geographic signatures. We made some preliminary experiments to assess new *combination metrics* for computing relevance based on geographic signatures.

The rest of this paper is organised as follows: Section 2 depicts our assembled GIR system, and describes in detail each module of our prototype. Section 3 presents our experiments and Section 4 analyses the results. Section 5 ends the paper with conclusions and directions for future work.

2 System Description

Figure 1 presents the architecture of the GIR system assembled for GeoCLEF 2007. The GeoCLEF topics are automatically parsed by QueOnde and converted into *<what, spatial relationship, where>* triplets. The QuerCol module performs term and geographic query expansion, producing query strings consisting of query terms and a query geographic signature (Q_{Sig}).

CLEF documents are loaded into a repository, becoming available to all modules. Faísca is a text mining module specially crafted to extract and disambiguate geonames, generating geographic signatures for each document (D_{Sig}). Sidra5, our index and ranking module, generates text indexes from the documents and geographic indexes from their geographic signatures. Sidra5 also receives the queries generated by QuerCol as input, and generates final GeoCLEF runs in the `trec_eval` format. All these components rely on a geographic ontology for geographic reasoning, created using our own geographic knowledge base, GKB [5].

2.1 Geographic Ontology

The geographic ontology is a central component of our GIR system, providing support for geographic reasoning for all modules. It models both geographic concepts and the relationships between concepts in

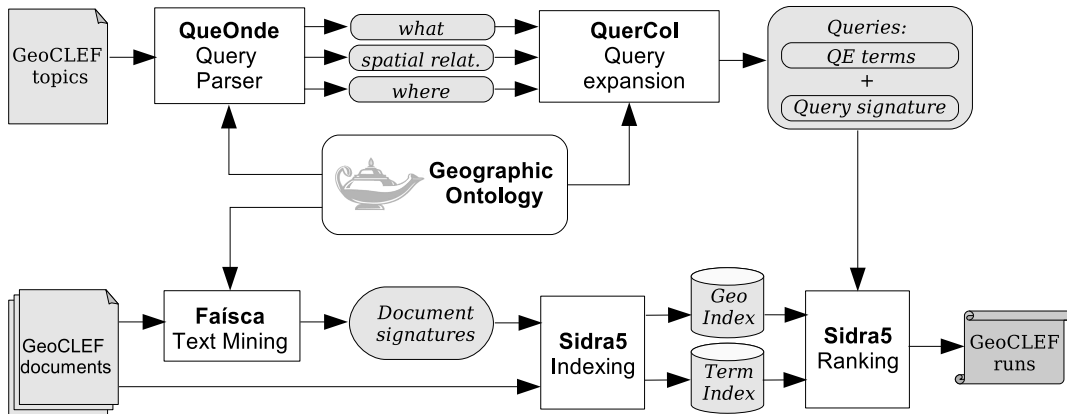


Figure 1: Architecture of the GIR system assembled for GeoCLEF 2007.

Physical Domain				Administrative Domain	
Island	205	Sea	5	Place	4023
Airport	107	Cathedral	3	ISO-3166-2	3976
River	86	Ocean	2	Administrative division	3212
Mountain	85	Mountain Range	2	Agglomeration	751
Lake	66	Strait	1	ISO-3166-1	239
Circuit	63	Channel	1	Capital city	233
Region	23	Planet	1	Total	12434
Continent	7	Total	657		
Names			14408	Centroids	4204
Features			13091	Bounding boxes	2083
Feature Types			21	<i>adjacent</i> relationships	11307
				<i>part-of</i> relationships	13762

Table 1: Statistics of the geographic ontology.

an hierarchical scheme. The geographic data come from several public sources, and include names for places and other geographic features, feature types, adjectives, relationships between concepts (*adjacent* and *part-of*), demographic data, spatial coordinates and bounding boxes [9].

The improvements made to the ontology for this year’s participation were twofold: i) update of the GKB conceptual model to directly support multilingual names for geographic references, and ii) the addition of new features that we found missing after inspecting the GeoCLEF topics for 2007. The GKB 2.0 model now supports relationships between feature types, a better property assignment for features and feature types, and a better control of information sources [6]. Most of the ontology enrichment was carried out in the physical domain, with the addition of new feature types like airports, circuits and mountains, along with their instances in the GKB. Table 1 presents the statistics of the ontology used in the evaluation.

2.2 QueOnde Query Parser and QuerCol Query Expansion

On the query side, we developed a new geographic query parsing module, QueOnde. The geographic query expansion module, QuerCol, introduced for last year’s participation [4, 9], was improved for also handling feature types and spatial relationships.

QueOnde automatically converts GeoCLEF topic titles into $\langle \textit{what}, \textit{spatial relationship}, \textit{where} \rangle$ triplets with the help of the geographic ontology and a set of manually-crafted context rules for capturing and disambiguating spatial relationships, features and feature types. QueOnde also participated on the 2007 GeoCLEF Query Parsing subtask [16].

The QuerCol module is able to expand the thematic (*what*) and the geographic (*where*) parts of a query separately. The *what* is expanded through blind relevance feedback (RF) [13], while the *where* is expanded

by a new algorithm, which decides the geographic expansion strategy to be performed based on features and feature types present on a query [3].

When feature types are present in the query, they may mean two things: i) the user is disambiguating the geoname, because it can be associated to other geographic concepts (e.g., *City of Budapest* and *Budapest Airport*); or ii) the user is designating a set of concepts as a scope of interest (e.g., *Airports of Hungary*). In case i), the feature type is disambiguating the geographic concept given by the feature *Budapest* as the scope of interest, while in case ii), the feature type is designating a group of geographic concepts of the scopes of interest, requiring additional geographic reasoning to obtain the corresponding concepts.

The geographic query expansion step of our GIR system is now guided according to the spatial relationship, features or feature types specified on the query. For instance, in the CLEF topic #74, *Ship traffic around Portuguese islands*, QuerCol considers *in* as the spatial relationship, *Portugal* as a feature name and *islands* as a feature type, and it reasons that the scope of interest is all geographic concepts of type *island* that are part of Portugal: *São Miguel, Santa Maria, Formigas, Terceira, Graciosa, São Jorge, Pico, Faial, Flores, Corvo, Madeira, Porto Santo, Desertas* and *Selvagens*.

2.3 Faísca

The text mining module Faísca parses the documents for geonames, generating geographic signatures for each document. Faísca relies on pattern matching from a gazetteer generated from the geographic ontology, containing all concepts represented by their names and respective feature types. Consider the following (fictional) example for the geoname *Lisbon*, which is associated to multiple geographic concepts in the ontology. The gazetteer would have the following pattern entries:

```
city $ Lisbon: 1
Lisbon city: 1
district $ Lisbon: 2
Lisbon district: 2
Street $ Lisbon: 3
Lisbon Street: 3
(...)
Lisbon: 1,2,3,(...)
```

The left size of these entries contains the text patterns to be matched, in [*<feature type>* \$ *<feature>*] and in [*<feature>* *<feature type>*] formats (being the former one more common for Portuguese texts, and the latter one for English texts), while on the right side there is an *identifier* of the corresponding geographic concept in the ontology. The character \$ means that an arbitrary term or group of terms is allowed to be present between the feature and the feature type, in order to avoid different stopword and adjective patterns. This approach immediately captures and grounds all geonames into their unique concept identifiers, without depending on hard-coded disambiguation rules. In the end, we have a *catch-all* pattern, which is used when the geoname found in the document does not contain any kind of external hints on its feature type. For these cases, we assign all identifiers of geographic concepts that are associated with the geoname *Lisbon*.

The geographic signatures (D_{Sig}) generated by Faísca consist on a list of concept identifiers and a corresponding *confidence measure* (*ConfMeas*) normalized to [0,1], that represents the confidence that the feature is part of the document scope. *ConfMeas* is obtained through an analysis of the surrounding concepts on each case, in a similar way as described by Li et al. [8]. Geonames on a text are considered as qualifying expressions of a geographic concept when a direct ontology relationship between the geonames is also observed. For example, the geoname *Adelaide* receives a higher *ConfMeas* value on the document signature if an ontologically related concept, such as *Australia*, is nearby on the text. If so, the feature *Australia* is not included in the D_{Sig} , because it is assumed that it was used to disambiguate *Australia*, the more specific concept. An excerpt of four document signatures (one per line) as generated by Faísca from the GeoCLEF collection is given below:

```
LA072694-0011: 5668[1.00]; 2230[0.33]; 4555[0.33]; 4556[0.33]; 4557[0.33]
LA072694-0012: 5388[1.00]; 5389[1.00]; 5390[1.00]; 12097[1.00]; 6653[0.67]
LA072694-0013: 369[1.00]; 225[0.33]; 452[0.33]; 7[0.33]; 367[0.33]; 137[0.33]
LA072694-0014: 6653[1.00]; 6654[1.00]; 347[1.00]
```

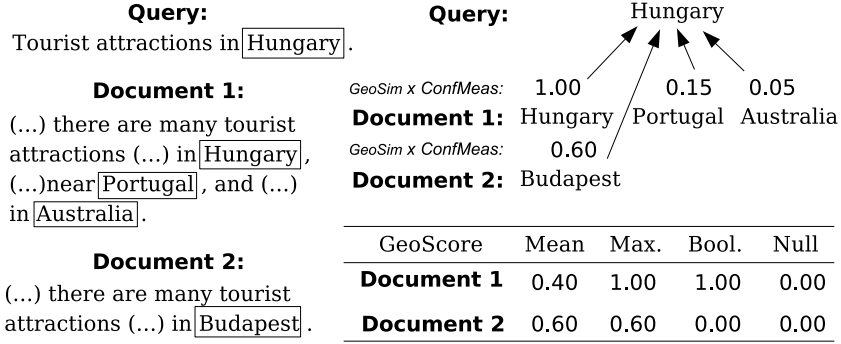


Figure 2: Example of the calculation of the four *GeoScore* combination metrics.

2.4 Sidra5

Sidra5 is a text indexing and ranking module with geographic capabilities based on Managing Gigabytes for Java (MG4J) [1]. It uses a standard inverted term index provided by MG4J, and a geographic forward index of $[docid, D_{Sig}]$ that maps the id of a document to the corresponding D_{Sig} generated by Faísca.

To retrieve documents, Sidra5 first uses the *what* part of the query and the term index to retrieve the top 1000 documents. Afterwards, the D_{Sig} of each document is retrieved with the help of the geographic index. Finally, the document score is obtained by combining the Okapi BM25 *text score* [12], normalized to $[0,1]$ (*NormBM25*) as defined by Song et al. [15], and a *geographic score* normalized to $[0,1]$ (*GeoScore*) with equal weights:

$$Ranking(query, doc) = 0.5 \times NormBM25(query, doc) + 0.5 \times GeoScore(query, doc) \quad (1)$$

The calculation of *GeoScore* begins with the computation of the geographic similarity *GeoSim* for each pair (s_1, s_2) , where s_1 in Q_{Sig} and s_2 in D_{Sig} , through a weighted sum of four heuristic measures (discussed in our 2006 GeocLEF participation [9]): Ontology (*OntSim*), Distance (*DistSim*), Adjacency (*AdjSim*) and Population (*PopSim*) similarity measures.

$$GeoSim(s_1, s_2) = 0.5 \times OntSim(s_1, s_2) + 0.2 \times DistSim(s_1, s_2) + 0.2 \times PopSim(s_1, s_2) + 0.1 \times AdjSim(s_1, s_2) \quad (2)$$

Having geographic signatures with multiple concepts requires adding aggregation metrics to *GeoScore* for handling the different *GeoSim* values that a $(query, doc)$ pair can generate. We experimented four metrics: Maximum, Mean, Boolean and Null.

Maximum: *GeoScore* is the maximum *GeoSim* value computed between a $(query, doc)$ pair.

$$GeoScore_{Maximum}(query, doc) = \max(GeoSim(s_1, s_2) \times ConfMeas(s_2)), s_1 \in Q_{sig} \wedge s_2 \in D_{sig} \quad (3)$$

Mean: *GeoScore* is the average *GeoSim* values computed between a $(query, doc)$ pair.

$$GeoScore_{Mean}(query, doc) = avg(GeoSim(s_1, s_2) \times ConfMeas(s_2)), s_1 \in Q_{sig} \wedge s_2 \in D_{sig} \quad (4)$$

Boolean: *GeoScore* equals 1 if there is a common concept in a $(query, doc)$ pair, and equals 0 otherwise.

$$GeoScore_{Boolean}(query, doc) = \begin{cases} 1 & \text{if } \exists s_1 = s_2, s_1 \in Q_{sig} \wedge s_2 \in D_{sig} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

Null: $GeoScore_{Null}$ is always 0, turning off the geographic scores. This is used as a baseline metric for comparing results obtained with the other metrics.

Run	Description
1	Geographic QE before RF. Classical text retrieval.
2	Geographic QE before RF, GIR with Mean geoscore.
3	Geographic QE before RF, Maximum geoscore.
4	Geographic QE after RF, Mean geoscore.
5	Geographic QE after RF, Maximum geoscore.

Table 2: Runs submitted to GeoCLEF 2007.

The computation of the four *GeoScore* metrics is illustrated in Figure 2, which presents a fictional query (*Hungary*), and two document surrogates, along with the $GeoSim \times ConfMeas$ values and final *GeoScore* values.

3 Runs

Table 2 summarises the submitted runs, a total of 10: five on the Portuguese monolingual subtask and five on the English monolingual subtask. Our runs aimed to:

- evaluate if the current GeoIR approach of treating geonames in a separate geographic ranking obtains better results than treating geonames as terms in a standard IR approach;
- determine which *GeoScore* combination metrics is best. We experimented the $GeoScore_{Mean}$ and $GeoScore_{Maximum}$ on our runs. The $GeoScore_{Boolean}$ and $GeoScore_{Null}$ metrics were later included in post-hoc experiments;
- measure the importance of the geographic query expansion before or after the relevance feedback step.

We generated initial queries from the topic titles to obtain *initial runs* for the RF. We used 32 top-k terms and 20 top-k documents as parameters for the blind relevance feedback [4]. The final query string combines expansion terms by aggregating semantically related concepts with the help of the MG4J logic operators, following the suggestions of Mitra et al. [11], and the concept identifiers from the Q_{Sig} .

The *Terms only* experiment (run 1) uses early geographic reasoning to generate a Q_{Sig} . Yet, it uses the names of geographic concepts as standard terms in the generation of the initial and final runs, meaning that this run uses only classical text retrieval.

The other runs use the text and geographic scores for ranking documents: *Geographic QE before RF* experiments (runs 2 and 3) considers the Q_{Sig} as the *where* part of the initial query, for initial run and final run generation, while the *Geographic QE after RF* experiments (runs 4 and 5) use only the captured concepts on the topic title as the *where* part for the initial run generation, and the Q_{Sig} on the final run generation. The *Terms/GIR* runs on these experiments differ by the use of the initial run generated in the *Terms only* experiment.

4 Results

Unfortunately, the runs submitted to GeoCLEF were hampered by programming errors in our GIR prototype, and so the obtained poor MAP values did not allow us to draw any early conclusions regarding our experiments. After some code revision, we managed to obtain more significant MAP values and conducted additional experiments with the fixed GIR prototype. The MAP values presented on Table 3 refer only to the post-hoc experiments.

We observed that the $GeoScore_{Mean}$ produces poor MAP values, because long document signatures tend to cause query drifting. $GeoScore_{Maximum}$ and $GeoScore_{Boolean}$ revealed to be much more robust, and the $GeoScore_{Boolean}$ metric has the best MAP values for Portuguese. This is explained in part because the $GeoScore_{Maximum}$ is highly dependent on the heuristics used, and these are dependent on the quality of the geographic signatures and the quality of the ontology, while the $GeoScore_{Boolean}$ metric is more straightforward on assigning maximum scores for geographically relevant documents. This difference also

	GeoScore	Terms only	Geo. QE before RF	Geo. QE after RF	Terms/GIR
Initial run		0.210	0.126	0.084	0.210
Final Run	Maximum		0.125	0.104	0.205
	Mean	0,233	0.022	0.021	0.048
	Boolean		0.135	0.125	0.268
	Null		0.115	0.093	0.021
a) Results for the Portuguese monolingual subtask.					
Initial run		0,175	0.086	0.089	0.175
Final Run	Maximum		0.093	0.104	0.218
	Mean	0.166	0.043	0.044	0.044
	Boolean		0.131	0.135	0.204
	Null		0.081	0.087	0.208
b) Results for the English monolingual subtask.					

Table 3: MAP results obtained for the post-hoc experiments.

means that there are more irrelevant documents that are being scored higher than relevant documents being scored lower by the $GeoScore_{Maximum}$.

Regarding the geographic query expansion before or after the RF, we found that early geographic expansion results in a better generation of initial runs (0.126 versus 0.084), meaning that more relevant documents are present on the top-k docs, thus improving the results from the RF step.

Using geonames as terms on the term index instead of geographic concepts still gets better results in the initial run (0.210 versus 0.126). The final run obtained without performing geographic ranking improves the MAP value to 0.233. We were intrigued with the consistent better results obtained with the *Terms Only* experiment. The good MAP value obtained by its initial run (0.210) suggested an experiment with this initial run, followed by a term and geographic expansion to generate a final query with a geographic signature, and ending in a GIR retrieval just like the other experiments. This *Terms/GIR* experiment obtained an MAP of 0.268 for the $GeoScore_{Boolean}$ metric, the highest MAP value of our post-hoc experiments.

Regarding the English experiments, we observe similar trends as in the Portuguese experiments. The slightly lower values are consequence of the quality of the ontology, which is more complete with Portuguese feature names. Also, we observe that the $GeoScore_{Maximum}$ outperformed the $GeoScore_{Boolean}$ geoscope values for the *Terms/GIR* experiment, which prompt us to make further analysis on the meaning of the observed differences between these two metrics.

5 Conclusions

This year’s participation was a deception in terms of results for official runs, but we accept it as the consequence of deciding to develop a totally renewed and untested GIR system. Yet, the post-hoc experiments drew some interesting results for understanding why the GIR approaches are still outperformed by classic IR approaches. Our *Terms/GIR* experiments manage to obtain the highest MAP values, which might shed some light on this problem and suggest that there may be more efficient ways to introduce geographic reasoning in a GIR system.

The approaches of this year and last year’s participations are both very dependent of the quality of the geographic ontology. 25% of the relevant documents contained geonames that were not in our ontology, and we found that we have poor results when handling queries with unknown geonames. In addition, the ontology is not comprehensive on coordinates and population data to serve the geographic heuristics. We need to make further experiments with a more complete ontology, in order to better evaluate the fitness of the geographic similarities.

We also believe that our results could be improved with a more robust term query expansion module, as the current query expansion through blind relevance feedback is basic and does not produce significant improvements. We are also aware that some of the blame may be on the query construction step, as the readaptation for the MG4J syntax was overlooked. Our post-hoc experiments used RF parameters of eight top-k terms and five top-k docs, and used different logic operations for query construction. These changes

resulted in significant improvement of the results, showing that we still have some tuning to do in the term query expansion step.

Finally, we conclude that this new GIR approach has its merits, and may be further improved to produce good results. Yet, it is still on its early steps, so our next work is to mature the approaches and develop a stable GIR prototype for further experiments.

Acknowledgements

We thank Joana Campos, for developing the text mining module, and Catarina Rodrigues for managing the geographic data. Our participation was jointly funded by the Portuguese government and the European Union (FEDER and FSE) under contract ref. POSC/339/1.3/C/NAC (Linguatca), and partially supported by grants SFRH/BD/29817/2006 and POSI/SRI/40193/2001 (GREASE) from FCT (Portugal), co-financed by POSI.

References

- [1] Paolo Boldi and Sebastiano Vigna. MG4J at TREC 2005. In *Proceedings of the 14th Text REtrieval Conference, TREC 2005*. NIST Special Publication SP 500-266, 2005. <http://mg4j.dsi.unimi.it>.
- [2] Nuno Cardoso, Bruno Martins, Leonardo Andrade, Marcirio Silveira Chaves, and Mário J. Silva. The XLDB Group at GeoCLEF 2005. In Carol Peters et al, editor, *Accessing Multilingual Information Repositories: 6th Workshop of the Cross-Language Evaluation Forum, CLEF 2005*, volume 4022 of LNCS, pages 997–1006. Springer, 2006.
- [3] Nuno Cardoso and Mário J. Silva. Query Expansion through Geographical Feature Types. In *4th Workshop on Geographic Information Retrieval, GIR 07 (held at CIKM'07)*, Lisbon, Portugal, 9th November 2007.
- [4] Nuno Cardoso, Mário J. Silva, and Bruno Martins. The University of Lisbon at CLEF 2006 Ad-Hoc Task. In Carol Peters, editor, *Cross Language Evaluation Forum: Working Notes for the CLEF 2006 Workshop*, Alicante, Spain, 20–22 September 2006.
- [5] M. S. Chaves, M. J. Silva, and B. Martins. A Geographic Knowledge Base for Semantic Web Applications. In C. A. Heuser, editor, *Proceedings of the 20th Brazilian Symposium on Databases*, pages 40–54, Uberlândia, Minas Gerais, Brazil, 3–7th October 2005.
- [6] Marcirio Silveira Chaves, Catarina Rodrigues, and Mário J. Silva. Data Model for Geographic Ontologies Generation. In Luís Carriço José Carlos Ramalho, João Correia Lopes, editor, *XATA2007 - XML: Aplicações e Tecnologias Associadas*, pages 47–58. Universidade do Minho, Fevereiro 2007.
- [7] William A. Gale, Kenneth W. Church, and David Yarowsky. One Sense per Discourse. In *HLT '91: Proceedings of the Workshop on Speech and Natural Language*, pages 233–237. ACL, 1992.
- [8] Yi Li, Alistair Moffat, Nicola Stokes, and Lawrence Cavedon. Exploring Probabilistic Toponym Resolution for Geographical Information Retrieval. In *Proceedings of the 3rd ACM Workshop on Geographical Information Retrieval, GIR'2006*, Seattle, Washington, USA, 10th August 2006.
- [9] Bruno Martins, Nuno Cardoso, Marcirio Chaves, Leonardo Andrade, and Mário J. Silva. The University of Lisbon at GeoCLEF 2006. Alicante, Spain, 20-22 September 2006. To be published by Springer.
- [10] Bruno Martins and Mário J. Silva. A Graph-Based Ranking Algorithm for Geo-referencing Documents. In *Proceedings of ICDM-05, the 5th IEEE International Conference on Data Mining*, Texas, USA, November 2005.
- [11] M. Mitra, A. Singhal, and C. Buckley. Improving Automatic Query Expansion. In *Proceedings of the 21st Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval*, pages 206–214. ACM Press, 1998.
- [12] Stephen E. Robertson, Steve Walker, Micheline Hancock-Beaulieu, Aaron Gull, and Marianna Lau. Okapi at TREC-3. In *Proceedings of TREC-3, the 3rd Text REtrieval Conference*, pages 21–30, 1992.
- [13] J. J. Rocchio Jr. Relevance Feedback in Information Retrieval. In Gerard Salton, editor, *The SMART Retrieval System: Experiments in Automatic Document Processing*, pages 313–323. Prentice-Hall, Englewood Cliffs, NJ, USA, 1971.
- [14] Mário J. Silva, Bruno Martins, Marcirio Chaves, Ana Paula Afonso, and Nuno Cardoso. Adding Geographic Scopes to Web Resources. *CEUS - Computers, Environment and Urban Systems*, 30:378–399, 2006.
- [15] Ruihua Song, Ji-Rong Wen, Shuming Shi, Guomao Xin, Tie-Yan Liu, Tao Qin, Jiyu Zhang, Xin Zheng, Guirong Xue, and Wei-Ying Ma. Microsoft Research Asia at the Web Track and TeraByte Track of TREC 2004. In *Proceedings of the 13th Text REtrieval Conference, TREC-04*, 2004.
- [16] Xing Xie. Query Parsing Task Proposal for GeoCLEF 2007. <http://ir.shef.ac.uk/geoclef/2007/Query-Parsing.htm>.