

The XLDB Group at the CLEF 2005 Ad-Hoc Task

Nuno Cardoso¹, Leonardo Andrade¹, Alberto Simões², and Mário J. Silva¹

¹ Grupo XLDB - Departamento de Informática

Faculdade de Ciências da Universidade de Lisboa

² Departamento de Informática, Universidade do Minho

{ncardoso, leonardo, mjs}@xldb.di.fc.ul.pt, ambs@di.uminho.pt

Abstract. This paper presents the participation of the XLDB Group in the CLEF 2005 ad-hoc monolingual and bilingual subtasks for Portuguese. We participated with an improved and extended configuration of the tumba! search engine software. We detail the new features and evaluate their performance.

1 Introduction

In 2004, the XLDB Group made its debut participation in CLEF, on the monolingual ad-hoc Portuguese retrieval task [1]. The main goals were to obtain hands-on experience in joint evaluations of information retrieval (IR) and evaluate tumba!, our web search engine [2] on this task. We learned that we had to come up with new approaches and methods, as the strategy for searching and indexing large web collections is different than when querying the kind of document collections used in the CLEF ad-hoc task.

This year, we embraced the ad-hoc task with the objective of evaluating new methods and algorithms for the task:

- Implementation of new logic operators on query strings to support expanded queries
- Development of new methods for using all the topic information provided and merging the combined result sets.
- Topic translation for submission of English to Portuguese bilingual runs.

This paper is organized as follows: Section 2 describes our system and enumerates the main changes from last year's configuration. In Section 3, we present our evaluation goals and submitted runs. Section 4 presents the results obtained. Section 5 summarises our conclusions.

2 Improvements

One of the main lessons learned from last year's CLEF ad-hoc task participation was that IR in large web collections is quite different from IR on small text collections. Simple adjustments to a web search engine aren't sufficient if we want to use all the information provided for each topic instead of just a few terms to query the CLEF ad-hoc collection. This motivated the development of a set of new software, to handle the task properly.

We developed a new query expansion module that generates alternative queries from the descriptions given. This module, called QuerCol (Queries Collator) is external to the core tumba! search engine, but has an essential role in the production of the runs we submitted to CLEF in 2005.

We also improved tumba! in its capability to properly rank poorly linked and tagged documents. We developed an algorithm based on $TF \times IDF$ weighting to rank the results for CLEF 2005, added support for the 'OR' operator in query strings, and implemented new result set merging algorithms.

With these new modules, our group is now taking the first steps to include the basic set of components required for serious participation on in this kind of IR task – robust stemming, weighting scheme and blind feedback [3].

In the remainder of this section, we detail the design of QuerCol, the newly developed query expansion module, and the improvements made to the query processing sub-system of tumba!.

2.1 Query Expansion

The main conclusion of our CLEF 2004 participation was that, in order to achieve higher recall values, we need to expand the title terms into alternative variants, as collections include many documents relevant to the query topic without all the topic terms [1]. So, this year we created multiple queries for each topic, based on synonyms, morphological and lexical expansion of the title terms, and a selection of other terms from the topic description.

Query strings can now include the 'OR' (disjunction) operator, which wasn't supported by the query server that we had in 2004. This enabled us to make extensive use of synonyms and morphological variations of the title terms. Other systems and former CLEF participants, like Nateau et al, experimented query expansion modules based on the 'OR' operator [4], and that inspired us to start QuerCol.

QuerCol generates queries from a given topic using the following approach:

1. *Eliminate common stop-words and CLEF-related stop-words.* The latter include terms like 'document' and 'relevant', which are frequent in topic descriptions. We obtain these by selecting the top 5 most frequent terms from all topics.
2. *Obtain title concepts.* After stop-word elimination, we assume that all remaining title words are root-terms of Boolean expressions in the disjunctive normal form, each representing a **concept**, which must be present in all query strings derived from the topic. We used *jspell* to expand morphologically the title concepts [5,6]. *Jspell* is a morphological analyser based on derivation: words are created applying a set of rules over a root term. This way, it is easy to check the root term and apply rules to create word derivations for each title concept. From these, we only pick those having a frequency of at least 5 in the collection.
3. *Obtain expanded concepts.* For each topic title, we take the terms as a conjunction query, which is submitted to the tumba! instance indexing the CLEF ad-hoc collection. Then, we measure the $TF \times IDF$ value for each term in the topic's set of words, for each document in the obtained result set. We rank the top 8 terms and discard those with a document frequency lower than 5 in the collection. The selected terms are called **expanded concepts**.

4. *Compute the similarity between the title concepts and the expanded concepts.* For instance, if the title concepts are *shark* and *attack*, and the term *strike* is selected as an expanded concept, we want to relate it to the *attack* concept, to create a query like *shark attack OR shark strike*. We used a database of term co-occurrences of Portuguese terms developed by the Porto node of Linguatca, built from two Portuguese corpora, CETEMPublico [7] and WPT 03 [8]. In the example above, we queried the database for the top-20 terms that co-occur after the term *shark*. If *strike* is in the result, we can say that the two terms belong to the same concept, and we add *strike* to the *attack* concept term list.

If an expanded concept isn't associated to a concept, it is later added to the query string as a disjunction. This means that expanded concepts don't influence the result set lists, but contribute to weighting the documents containing them.

5. *Query string generation.* In the end, each title concept is defined as a list of terms, selected both from the expanded concepts and from the morphological expansions of the initial title terms. With all the lists of concepts for each topic, we compute all term combinations as a $m \times n$ matrix of m concepts $\times n$ term list size for each concept, and finally we merge them with disjunction operators to generate a single query string.

For the English to Portuguese ad-hoc bilingual subtask, we devised the two following approaches:

1. Using the Babelfish web translation tool (<http://babelfish.altavista.com>). The topic strings were extracted and sent one at a time to the translator and the translations replaced the original topic strings.
2. Using Example Based Machine Translation (EBMT) methods in parallel corpora [9]. The translations were made from a translation memory built from multilingual thesauri freely available on the Internet (EuroVoc, Unesco thesaurus and others). The thesauri have not only simple term entries but also multi-word entries that help in the translation of some word sequences. The translation memory was then used to translate word sequences of the topics file. Words without a corresponding entry in the translation memory were individually translated using Babelfish.

2.2 Weighting and Ranking

Sidra is the indexing and ranking system used in the tumba! web search engine [10]. Sidra provides support for "all the terms" searches, exact phrase queries and field searches that restrict result sets to a specific subdomain or document format. Sidra was primarily designed to rank web documents, as its original ranking function relied mainly on metadata such as links' anchor text, URL strings and page titles. However, it performs poorly when handling document collections with scarce metadata, such as the CLEF ad-hoc collection. Sidra does not perform term stemming; the index terms are all the single words, indexed as a full inverted file.

To improve the performance of Sidra on CLEF, we made two major enhancements:

1. Implement a weighting algorithm based on $TF \times IDF$. This enables us to tackle the absence of meta-data, and to have a baseline for a future implementation of a full Okapi BM 25 schema [11].

2. Develop support for disjunction of terms. Query strings submitted to Sidra may now include the 'OR' and 'AND' logic operators, as long as the query respects the Disjunctive Normal Form.

As Sidra query servers handle each conjunction as a simple query, support for the 'OR' operator consisted in devising strategies for merging the result sets of ranked documents obtained in each sub-query. We used two simple approaches:

Weight Merge: The final result set is obtained by sorting the weights of each result on the combined result set. The final weight of a document present in more than one result set is the sum of the weights of the document in each result set.

Round-Robin Merge: The final result set is generated by sorting the result sets by the weight of the top ranked document in the result set. Then, documents are picked from each result set using a round-robin rule. Documents already picked to the merged result set are ignored.

3 Runs

For the ad-hoc task, we submitted 5 runs for the Portuguese ad-hoc monolingual subtask (4 regular runs plus one mandatory run) and 4 for the English to Portuguese ad-hoc bilingual ad-hoc subtask. As we were testing implementations of the 'OR' operator on *tumba!*, we selected the result set merging methods as a parameter to measure which produced better results. Hence, we applied the Weight Merge algorithm to half the runs plus the mandatory run, and Round Robin Merge to the other half (see Table 1).

Table 1. Runs submitted to the ad-hoc task

Monolingual				
Query	Manual		Automatic	
Fusion	Weight	Round Robin	Weight	Round Robin
Run	XLDBTumba01	XLDBTumba05	XLDBTumba02 XLDBTumba09	XLDBTumba06

Bilingual				
Query	EBMT translation		Babelfish Translation	
Fusion	Weight	Round Robin	Weight	Round Robin
Run	XLDBTumba03	XLDBTumba07	XLDBTumba04	XLDBTumba08

In the monolingual subtask, we created runs XLDBTumba01 and XLDBTumba05 by manually adding all kinds of synonyms and morphological expansions that seem reasonable to the queries. We used it as a baseline for evaluation against other submitted runs. For runs XLDBTumba02 and XLDBTumba06, QuerCol automatically generated the queries. We aimed at obtaining result sets of the same level of quality as for manually created runs, as QuerCol used the same query creation approach. XLDBTumba09 is a mandatory run, with query strings automatically generated from the topics' title and description fields only.

On the bilingual subtask, the goal of our participation was to have a preliminary evaluation of the EBMT systems being developed at the Braga node of Linguateca.

4 Results

Figure 1 and Table 2 show the obtained results. One of our main goals was to compare the two result sets merging strategies, and in the end the Weight merge method outperformed the Round-Robin method. A deeper analysis on the results will provide valuable hints on the result set merging mechanism to implement for disjunctive queries.

Manual query creation (runs 01 and 05) performed better than automatic query creation (runs 02 and 06). Further analysis on the obtained results will also provide good hints for improving QuerCol to narrow the difference.

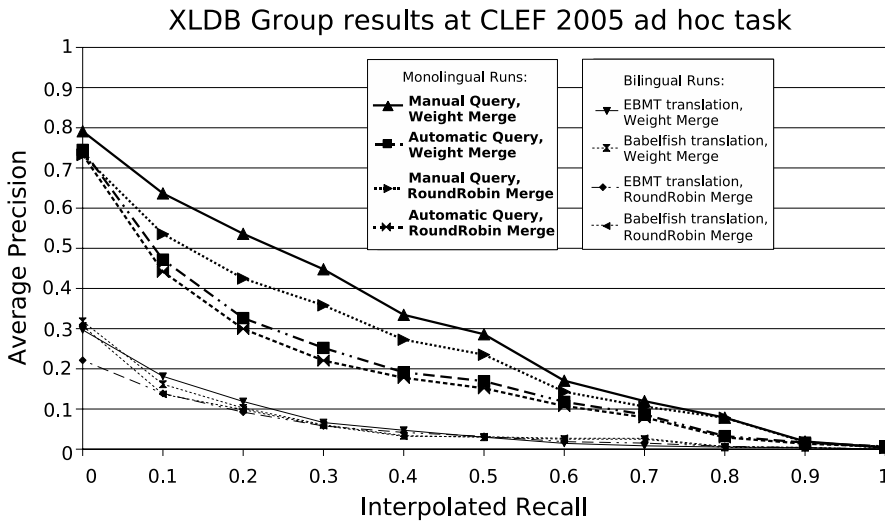


Fig. 1. Results of the XLDB Group on ad-hoc monolingual (thick lines) and bilingual subtasks (thin lines)

Table 2. Overall results on all runs

Run label	Retrieved	Relevant	Ret_rel	Avg. Prec.	R-Prec.	Overall Prec.	Overall Recall
XLDBTumba01	12595	2904	1675	29.0%	34.3%	13.3%	57.7%
XLDBTumba02	5546	2904	986	19.7%	23.2%	17.8%	34.0%
XLDBTumba05	12595	2904	1666	24.0 %	30.6%	13.2%	57.4%
XLDBTumba06	5546	2904	985	18.1%	22.5%	17.8%	34.0%
XLDBTumba03	4875	1991	605	5.8%	8.0%	12.4%	30.4%
XLDBTumba04	6774	2156	299	5.5%	7.4%	4.4%	13.9%
XLDBTumba07	4875	1991	617	4.7%	7.2%	12.6%	31.0%
XLDBTumba08	6774	2156	301	5.3%	7.4%	4.4%	14.0%
XLDBTumba09	6521	2904	989	19.4%	22.9%	15.2%	34.0%

The results of the monolingual runs are much better than the bilingual. This is likely to be a consequence of some poor translations. We concluded that we were using thesauri with less quality than expected. As we have overlaps (alternative translations coming from different thesauri), some of the used translations came from the wrong thesaurus and were the source of the bad translation results. Table 2 shows that the runs using EMBT translation obtained more relevant results with less retrieved documents, which is an encouraging result.

The relative performance of the best of our runs compared to other groups' submissions is close to the median. There are a few queries where our performance is much worse than the median for reasons that we have yet to find. However, given that in 2005 our weighting algorithm was very simple, we believe that an improvement here would likely raise the performance level of our software in future evaluations.

5 Conclusion

The results we obtained this year show a major improvement over last year. This comes as a direct consequence of the changes made to our IR system. Some of the developments for this CLEF task will be incorporated in the next version of tumba!

We have also identified further improvements, like extending QuerCol with a Portuguese stemmer. This would create better term expansions and improve the 'clustering' of terms from the same concept. QuerCol's generated queries also revealed some flaws that we need to amend, as there are concepts with more than one term that shouldn't be handled separately (for instance, *Bill Clinton*). Some morphological expansions of title terms might also produce misleading variations. Finally, we could also incorporate the software developed for our participation in GeoCLEF 2005 to expand geographic names in queries [12].

Acknowledgements

We would like to thank Daniel Gomes, who managed the tumba! repository instance created to support our participation in this joint evaluation. Thanks also to the developers of the tumba! search engine and the Portuguese language tools used to assemble the runs. Our participation was partly financed by the Portuguese Fundação para a Ciência e Tecnologia through grants POSI / PLP / 43931 / 2001 (Linguateca) and POSI / SRI / 40193 / 2001 (GREASE).

References

1. Cardoso, N., Silva, M.J., Costa, M.: The XLDB Group at CLEF 2004. In Peters, C., ed.: Working Notes for the CLEF 2004 Workshop, Bath, UK (2004)
2. Silva, M.J.: The Case for a Portuguese Web Search Engine. In: Proceedings of ICWI-03, the 2003 IADIS International Conference on WWW/Internet, Algarve, Portugal, IADIS (2003) 411–418
3. Braschler, M., Peters, C.: 1-2. In: Cross-Language Evaluation Forum: Objectives, Results, Achievements. Volume 7. Kluwer Academic Publishers (2004) 7–31

4. Nateau, D., Jarmasz, M., Barrière, C., Foster, G., St-Jacques, C.: Using COTS Search Engine and Custom Query Strategies at CLEF. In C.Peters, ed.: Working Notes for the CLEF 2004 Workshop, Bath, UK (2004)
5. Almeida, J.J., Pinto, U.: Jspell – a Module for Generic Natural Language Lexical Analysis. In: Actas do X Encontro da Associação Portuguesa de Linguística, Évora (1994) 1–15 in Portuguese. <http://www.di.uminho.pt/~jj/pln/jspell11.ps.gz>.
6. Simões, A.M., Almeida, J.J.: Jspell.pm – a Morphological Analysis Module for Natural Language Processing. In: Actas do XVII Encontro da Associação Portuguesa de Linguística, Lisbon (2001) 485–495 In Portuguese.
7. Rocha, P., Santos, D.: CETEMPúblico: Um corpus de grandes dimensões de linguagem jornalística portuguesa. In: Actas do V Encontro para o processamento computacional da língua portuguesa escrita e falada (PROPOR'2000), ("Atibaia, São Paulo, Brasil)
8. Martins, B., Silva, M.J.: A Statistical Study of the WPT 03 Corpus. Technical Report DI/FCUL TR-04-1, Departamento de Informática da Faculdade de Ciências da Universidade de Lisboa (2004)
9. Somers, H.: Review article: Example based machine translation. *Machine Translation* **14** (1999) 113–157
10. Costa, M.: Sidra: a flexible web search system. Master's thesis, Faculdade de Ciências da Universidade de Lisboa (2004)
11. Robertson, S.E., Walker, S., Jones, S., Hancock-Beaulieu, M.M.: Okapi at TREC-3. In Harman, D.K., ed.: IST Special Publication 500-225: Overview of the Third Text REtrieval Conference (TREC 3), Gaithersburg, MD, USA, Department of Commerce, National Institute of Standards and Technology (1995) 109–126
12. Cardoso, N., Martins, B., Chaves, M., Andrade, L., Silva, M.J.: The XLDB Group at Geo-CLEF 2005. In Peters, C., ed.: Working Notes for the CLEF 2005 Workshop, Wien, Austria (2005)