UB at CLEF2004: Part 2 – Cross Language Medical Image Retrieval

Miguel E. Ruiz   Munirathnam Srikanth
State University of New York at Buffalo
{*meruiz,srikanth*}*@buffalo.edu*

Abstract

This paper presents the results of the University at Buffalo in CLEF 2004. Our efforts concentrated in two main tasks: Multilingual task using English as the initial query language, and medical image retrieval. The paper has been divided into two parts that will appear in the respective areas of the proceedings of workshop. Our Adhoc retrieval work used the TAPIR toolkit developed in house by the second author. Our approach focused on the validation and adaptation of the language model system to work in a multilingual environment and in exploring ways to merge results from multiple collections into a single list of results. The second part of the paper describes our experiments in multilingual image retrieval. Our work in image retrieval explores automatic query expansion using pseudo relevance feedback on the case descriptions to improve ranking of images retrieved by a CBIR system. Our results are quite good and show significant improvements with respect to our baseline.

1 Introduction

For CLEF 2004 we participated in the adhoc mono and multilingual retrieval as well as in the medical image retrieval. The goal of our participation in the adhoc retrieval task is to explore language modeling approaches for retrieval from non-English collections using the TAPIR (Text Analysis and Processing for Information Retrieval) toolkit which was originally developed for English and later modified to support ISO-Latin-1 encoding and Porter stemmers for European languages. Part one of this paper will present in detail our results for the monolingual retrieval task in French, Finnish and Russian, as well as the multilingual task using English queries to retrieve information in English, Finnish, French and Russian collections.

The second part of this paper presents our results for the medical image retrieval task. In this track our goal is to improve image retrieval by using retrieval feedback on the related case descriptions to re-rank the images retrieved by a CBIR system. Because our Language model system did not support retrieval feedback (which is a feature that was still under development by the time we worked on this task) we decided to use a version of the SMART retrieval system that we used in our participation in CLEF2003 [2].

2 Medical Image Retrieval

In the medical image retrieval task the query is a medical image and the goal is to generate a ranked list of images that are related to the original query. Our goal in this task was to explore ways to expand the initial image retrieval with the multilingual text of the case descriptions associated to each image. For this purpose we used a pseudo relevance feedback mechanism. The first step consists in performing retrieval using the database of images indexed with the Viper system. The top n images are used to locate the corresponding case descriptions. These case descriptions are used to build a query that is submitted to the text retrieval system to obtain other related case descriptions.

2.1 Collection Preparation and Indexing

The collection consists of $8,726$ images and $2,081$ cases descriptions that contain clinical information. Our initial inspection of the data revealed that there were 209 cases that have images associated with them but no textual information. We discarded these cases from our experiments because they would not be suited for our evaluation. In consequence, our text collection consists of $1,872$ cases.

We used the list of retrieved images by *Viper*, which was supplied by the organizers of this track. For this reason, our efforts in pre-processing concentrated in manipulating the text descriptions associated with these images.

We decided to use almost all tags included in the documents with the exception of dates, URLs, and personal information from the patients (i.e. birth date, age, etc). These tags were classified and grouped into 9 types:

- Textual description: this includes fields such as title, description, commentary, questions, and answers
- Diagnosis: The actual diagnosis associated to each case.
- Keywords and codes: This type includes keywords assigned to the case and radiology classification codes (ACR)
- Authors and organizations: Author, reviewer, hospital, department
- Language
- Orthopedic information: This includes all tags related to orthopedic annotations.
- Images: We added the list of image ids associated with each case.

Each of these types of information has its own characteristics that merit a different treatment during text processing and indexing. Our initial design creates a separate index for each type. The final score for ranking the retrieved cases is a weighted linear combination of each index score. Ideally, the weight of the contribution of each type should be determined experimentally. However, because we did not have a reliable way to estimate the contribution of each type to the final score of the document we decided to use the same weight for all parts.

Most of the case descriptions have a language field that indicates the language used in them. However, some case descriptions have text in both languages (French and English) and we were not sure how often this occurred in the collection. We use a simple algorithm to estimate whether the actual language used in the document corresponds with the language assigned to the case in the language field. This algorithm identifies stop words in English and French and computes a score for each language based on the proportion of English and French stopwords present in the document. Through this process we found that $1,693$ cases were in French, 177 were in English and 16 cases have text in both languages (i.e. French description with English comments). Given the nature of these bilingual texts we decided to build a single retrieval index for all documents instead of separating them into two sub-collections.

Our previous experience with medical documents has shown that using an aggressive stemming such as Porter's stemmer could reduce terms to roots that are actually quite different from their intended meaning. For example, "organization" is stemmed to "organ", which has a very different meaning from the original word. For this reason we use a simple stemming strategy that takes care only of plurals (in both English and French). We also used a stopword list that combined English and French stopwords and was manually reviewed to assure that it did not contain stop words that could have medical meaning (for example, the original stopword list from SMART includes

Weight	Term	Weight	Term
0.28935	im10654	0.08574	iliite
0.27833	im10361	0.08092	pied
0.26395	im11040	0.07697	acr33.3320
0.26294	im11114	0.07697	acr44.3320
0.25794	im10945	0.07697	im10362
0.25652	im10170	0.07212	l'èvolution
0.25585	im9832	0.07180	dèmasquage
0.25585	im9833	0.06919	sènile
0.25585	im9835	0.06819	kindyni
0.23936	im10916	0.06573	psoriasi
0.13769	sacro	0.06572	patiente
0.11502	bassin	0.06379	toutefoi
0.09297	iliaque	0.06258	im11042
0.08798	acr44.562	0.06258	im11041
0.08798	im10655	0.06169	collection

Table 1: Top 30 terms generated by the query expansion method for the first image query

"B" and "E" as a stop words, but if we discard this words it would be difficult for the system to distinguish between articles that talk about "vitamin B" and "vitamin E").

Indexing was performed using a version of the SMART system adapted to handle the ISO-latin-1 encoding. The documents were indexed using *atc* weighting (augmented term frequency, idf, and cosine normalization) while the queries used *atn* weighting (augmented term frequency, idf, no normalization).

2.2 Query Expansion

Our retrieval approach follows a classical pseudo relevance feedback method. The initial image is send as a query to *Viper* and the top ten images retrieved are used to build a query for the textual database. Our initial text query consists of the image ids of the top ten images retrieved (Note that we have added the list of image ids related to each case). We perform an initial retrieval step using these queries and retrieve the top 1000 cases. The top n cases are marked as relevant and will be used to obtain terms to expand the original query. The query expansion step uses Rocchio's formula to compute the weight of each of the terms as follows:

$$Q_{new} = \alpha \times Q_{orig} + \beta \times \frac{\sum_{D \in Rel} D}{R} + \gamma \times \frac{\sum_{D \ni Rel} D}{N - R} \qquad (1)$$

Terms are ranked according to the Rocchio's score and the top m terms are selected for expansion. We tried several values for the number of cases assumed to be relevant after the initial retrieval (n = 5, 10, 20) and for the number of terms used to build the expanded query (m = 20, 50, 100). Since we were not sure whether the usage of the original image ids would be important or not to the final retrieval we decided to use two different values for the coefficient α: 0 (don't take into account these original terms) and 1. The second coefficient (β) of the Rocchio's formula controls the contribution of the relevant documents. We set it to 64 because this is the most important information that will allow us to expand the query. The third coefficient γ controls the penalty assigned to terms that appear in the "non-relevant" documents (bottom 100 cases retrieved in the initial retrieval) and was set to 16. In summary, we tried two different sets of coefficients for the Rocchio expansion formula ($\alpha = 1$, $\beta = 64$ and $\gamma = 16$) and ($\alpha = 0$, $\beta = 64$ and $\gamma = 16$). An example of the expanded query is shown in Table 1.

The expanded query is then submitted to text retrieval system and the score of each retrieved case is assigned to the images associated with it. A final score for each image was computed by

	AvgP	above median
UBMedImTxt01	0.3488	26
UBMedImTxt02	0.3309	20
UBMedImTxt03	0.3291	20

Table 2: Performance of official runs in Medical Image Retrieval

combining the scores obtained from the image retrieval system and the text retrieval system. We use a linear combination of the scores to compute the final image score:

$$W_k = \lambda Iscore_k + \delta Tscore_k \tag{2}$$

where $Iscore_k$ and $Tscore_k$ are the scores assigned to the image k by the image retrieval system and text retrieval system respectively, λ and δ are coefficients that weight the contribution of each score. Usually the coefficients are estimated from experimental results. However, due to the lack of training data we decided tu use $\lambda = \delta = 1$ (observe that this simple addition of scores is possible due to the fact that both scores are scaled to be between 0 and 1).

3 Analysis of Results

We decided to submit three runs. The first run (UBMedImTxt01) used the top 10 documents to expand the query with the top 100 terms ranked by Rocchio's formula with coefficients $\alpha = 1$, $\beta = 64$ and $\gamma = 16$. This is a run that uses an aggressive expansion strategy and takes into account the image ids of the top ten images retrieved by *Viper* as actual terms. The second run (UBMedImTxt02) differs from the first run in the fact that the coefficient $\alpha = 0$ disregards the image ids as actual query terms. The third run (UBMedImTxt03) uses a more conservative strategy for expansion with only the top 5 cases and coefficients $\alpha = 1$, $\beta = 64$ and $\gamma = 16$.

Our official results are presented in Table 2. We were very pleased with the performance of the system since it shows a positive impact in improving relevance of the images retrieved. The best run UBMedTxt01 performed above the median in all queries and obtained the best performance of all official runs in automatic query construction. Our second run (UBMedImTxt02) performs 5.4% below our best run and performs above the median in 20 queries. Observe that the only difference between these two runs is that we use the ids of the images as actual terms for query expansion. These image ids work as anchors that reinforce the notion that cases that those cases, which have images associated with the assumed top 10 retrieved images, are regarded as relevant in our initial retrieval. Our third run (UBMedImTxt03) performs 6% below the best run and performs above the median systems in 20 queries. This third run uses a more conservative query expansion assuming that only the top 5 retrieved cases are relevant and perform query expansion.

We have to note that improvements to the final performance of the expanded queries are highly dependent of the quality of the initial set of images retrieved by the CBIR system. We asked two physicians (an specialist in pneumonology and an urologist) to help us validate the results of the retrieved images (although we did not use this feedback to change the ranking of the images or the way the system processed the queries). We asked them to give general feedback to understand whether the results retrieved by the system would make sense to a medical professional. This helped us to realize that some of the aspects of how a medical professional could use this type of system in their daily work. They also emphasized that the actual diagnostic is usually a complex process that includes not only the review of images but also the analysis of the clinical data that in many cases is more indicative of a specific diagnosis than the image itself. This seems to be corroborated by the fact that adding the text description of the actual cases associated to the image makes a significant difference.

Observe that because we have indexed French and English documents as a single collection the expanded query actually includes terms in both languages. A different approach could be to perform query expansion in two separate collections and then merge the results I n a single list.

Another approach could identify the language of a terms and add the corresponding translation. However, this will require the use of a specialized bilingual lexicon.

4 Conclusions and Future Work

Although we still have some more work to do regarding the interpretation of the results and the impact of combining text and image retrieval pseudo relevance feedback, we can conclude that this technique has a positive impact on the final performance for ranking relevant images.

Our method for preprocessing the actual structure of the cases have to be refined but it seems to work well for retrieval purposes.

We plan to add some extra query expansion using the UMLS Metathesaurus produced by NLM to add related medical phrases to the cases and verify whether this would actually improve performance.

5 Acknowledgement

We would like to thank Doctors Freddy and Adriana Karaush for their help in validating the results retrieved by the system and giving positive feedback that we will use for making general improvements to our system.

References

[1] Rocchio, J. J. Relevance feedback in information retrieval. In G. Salton (Ed.) The SMART Retrieval System: Experiments in Automatic Document Processing (pp.313.323). Englewood Cliff, NJ, Prentice Hall, 1971.

[2] Ruiz, M. E. Automatically generated phrases and relevance feedback for improving Cross-Language Information Retrieval. In Proceedings of the Cross Language Evaluation Forum (CLEF) 2003, August 21st . 22nd, Trondheim, Norway, 2003.

[3] Salton, G. (Ed.) The SMART Retrieval System: Experiments in Automatic Document Processing. Englewood Cliff, NJ, Prentice Hall, 1971.