MIRACLE at ImageCLEFphoto 2007:
Evaluation of Merging Strategies for Multilingual and Multimedia Information Retrieval

Julio Villena-Román[1,3], Sara Lana-Serrano[2,3],
José Luis Martínez-Fernández[1,3], José Carlos González-Cristóbal[2,3]

[1] Universidad Carlos III de Madrid
[2] Universidad Politécnica de Madrid
[3] DAEDALUS – Data, Decisions and Language, S.A.

jvillena@it.uc3m.es, slana@diatel.upm.es,
joseluis.martinez@uc3m.es, josecarlos.gonzalez@upm.es

Abstract

This paper describes the participation of MIRACLE research consortium at the ImageCLEF Photographic Retrieval task of ImageCLEF 2007. For this campaign, the main purpose of our experiments was to thoroughly study different merging strategies, i.e. methods of combination of textual and visual retrieval techniques. While we have applied all the well known techniques which we had already used in previous campaigns, for both textual and visual components of the system, our research has primarily focused on the idea of performing all possible combinations of those techniques in order to evaluate which ones may offer the best results and analyze if the combined results may improve (in terms of MAP) the individual ones.

The system includes three main modules. On one hand, apart from the search engine (Xapian or Lucene), the textual retrieval module includes parsers, stemming, stopword filtering, proper noun detection and semantic expansion components. On the other hand, the visual retrieval module is based on two well-known content-based engines: GIFT and FIRE. Finally, the merging module allows to use different operators (AND, OR, LEFT, RIGHT) to combine the outputs of the two previous subsystems and to calculate the result relevance based on different metrics (max, min, avg, max-min).

We finally submitted 110 multilingual textual (text-based) runs, 22 visual (content-based) runs and 21 mixed runs. Results in general show a poor performance for all groups, due to the characteristics of the image collection and the difficulty of the defined topics. The most interesting conclusion is that the defined merging strategies are successful as our best mixed experiment outperforms both the textual and visual experiments in which it is based, using the LEFT operator for the combination along with the max-min metric for computing the relevance.

Categories and Subject Descriptors

H.3 [Information Storage and Retrieval]: H.3.1 Content Analysis and Indexing; H.3.2 Information Storage; H.3.3 Information Search and Retrieval; H.3.4 Systems and Software; H.3.7 Digital libraries. **H.2 [Database Management]**: H.2.5 Heterogeneous Databases; **E.2 [Data Storage Representations]**.

Keywords

Linguistic Engineering, Information Retrieval, Image Retrieval, indexing, domain-specific vocabulary, Semantic Expansion, WordNet, Word Sense Disambiguation

1. Introduction

MIRACLE is a research consortium formed by research groups of three different universities in Madrid (Universidad Politécnica de Madrid, Universidad Autónoma de Madrid and Universidad Carlos III de Madrid) along with DAEDALUS, a small/medium size enterprise (SME) founded in 1998 as a spin-off of two of these groups and a leading company in the field of linguistic technologies in Spain. MIRACLE has taken part in CLEF since 2003 in many different tracks and tasks, including the main bilingual, monolingual and cross lingual tasks as well as ImageCLEF, Question Answering, WebCLEF and GeoCLEF tracks.

This paper describes our participation at the ImageCLEF Photographic Retrieval task of ImageCLEF 2007. Briefly, the goal of this task (fully described in [8]) is: given a multilingual statement describing a user specific information need, find as many relevant images as possible from the given multilingual document collections containing images as well as text. The reference database for this campaign is the IAPR TC-12 Benchmark [9], created under Technical Committee 12 (TC-12) of the International Association of Pattern Recognition (IAPR [10]). This collection contains 20,000 photos (mainly colour photographs) taken from locations around the world and comprises a varying cross-section of still natural images, annotated with semi-structured captions in English and German.

Participants are provided with a list of 60 topics which include a short textual title representing the research goal in 15 different languages (Danish, Dutch, English, Finnish, French, German, Italian, Japanese, Norwegian, Polish, Portuguese, Russian, Spanish, Swedish, and Simplified and Traditional Chinese) and, in addition, three image examples for each topic. The objective is to retrieve as many relevant images as possible from the given visual and multilingual topics.

For this campaign, the main purpose of our experiments was to thoroughly study the different merging strategies, i.e. methods of combination of visual and textual techniques. While we have applied all the well known techniques which we had already used in previous campaigns [11] [12], for both textual and visual components of the system, our research has focused on the idea to perform all possible combinations of those techniques in order to evaluate which ones may offer the best results.

All experiments are fully automatic, thus avoiding any manual intervention. None of the experiments incorporates relevance feedback, due to time constraints. We finally submitted 110 multilingual textual (text-based) runs, 22 visual (content-based) runs and 21 mixed runs (using a combination of both), elaborated on in the following sections.

2. System Description

Based on our experience in previous campaigns, to be able to execute a large number of runs that exhaustively cover all the combinations of the different techniques, we developed a flexible system, composed of a set of small components which can be easily added in different configurations and are executed sequentially to build the final result set.

Figure 1 presents an overview of the system architecture.
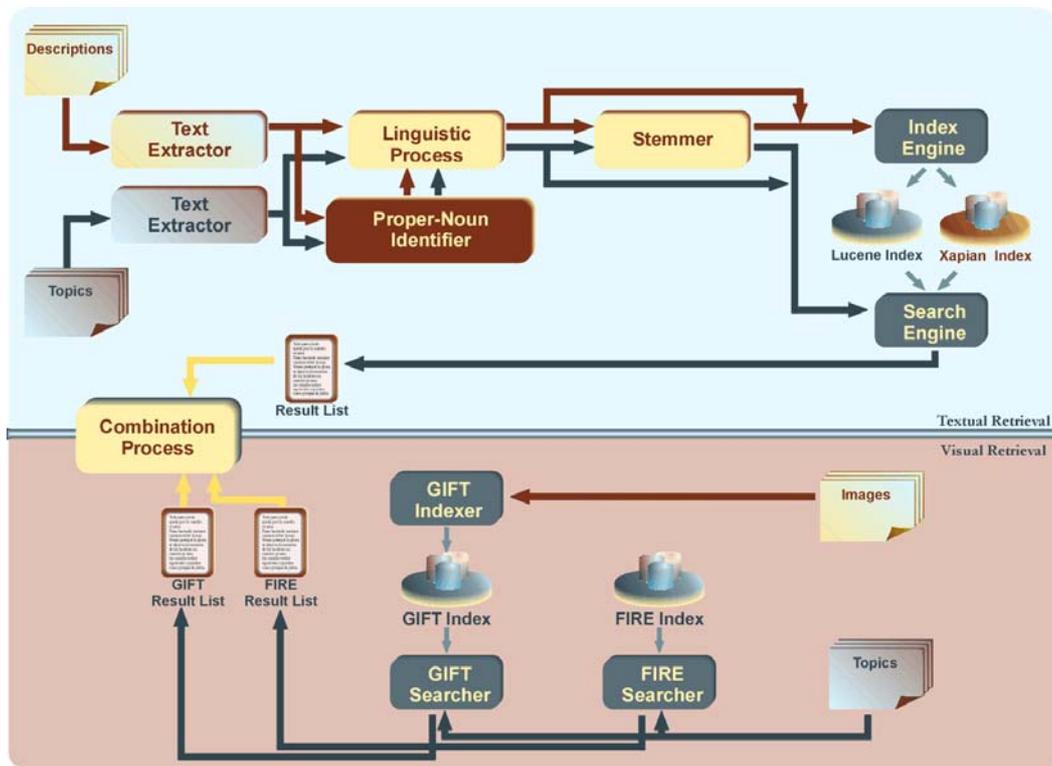
Figure 1. Overview of the system.

Our approach is based on the multidisciplinary combination of the text-based retrieval techniques using Lucene/Xapian with the GIFT/FIRE content-based retrieval. Thus, the system is logically built up from three different components:

- the **textual** (text-based) **retrieval module**, which indexes the IAPR TC-12 image descriptions to look for those descriptions that are more relevant to the text of the topic;

- the **visual** (content-based) **retrieval module**, that provides the IAPR TC-12 images which are more similar to the given topic images;

- and, finally, the **merging module**, which uses different operators to combine the outputs of the two previous subsystems to provide the final results.

2.1. Textual Retrieval

Since MIRACLE has taken part in ImageCLEF (or, in general, in CLEF), different linguistic and statistical techniques have been developed to be used in the text-based part of the different tasks [11] [12]. This year, the main goal was to make an exhaustive study of these diverse methods, combining them in all possible ways and testing the results achieved. A flexible and configurable system has been developed for this purpose.

The list of components includes:

- **Proper Noun Detection**. A heuristic based module to detect the appearance of entities in the text, based on a finite state automaton.

- **Linguistic Analyzer**. A component to obtain morphosyntactic analysis (POS tagging) and lemmatization. The target language considered in MIRACLE experiments has been English (although the IAPR TC-12 collection was also available in German and Spanish) so the Charniak parser [2] has been integrated. The combination of this parser with (Euro) Wordnet [5] has been used to provide lemmatization.

- **Stemming**. Of course, a component to obtain stems for words has been included as part of the indexing process. The implementation considered for this component has been the Porter's stemmer [14].

- **Stopword Detection**. One of the usual steps to take into consideration in the indexing process is the exclusion of semantic empty words, also called stopwords [3] [15].

- **Semantic Expansion**. Words appearing in the query can be expanded using Wordnet [5] and an implementation of a Word Sense Disambiguation algorithm [13]. Two working modes have been considered, one where each word is expanded with all the words considered as synonyms by Wordnet, and another one where a disambiguation process is carried out to try to filter out synonyms for the wrong senses.

- **Search Engine**. Two different search engines have been used, Lucene [1] and Xapian [16].

The former methods have been combined in all possible ways and applied to the text extracted from the fields (either individually or combined) in which image descriptions are divided. As described in [9], image captions include the following fields: Title, Description, Notes, Location, Date, Image and Thumbnail. However, the Description field is left out in the 2007 collection.

The final goal of this system configuration was to evaluate which combination of techniques was the best, considering last year relevance assessments. However, there were many possible combinations (although some of them were illogic, like applying semantic expansion to stems). Therefore only those with the best results on 2006 data where selected.

2.2. Visual Retrieval

For this part of the system, we resorted to two publicly and freely available Content-Based Information Retrieval systems: GIFT (GNU Image Finding Tool) [7] and FIRE (Flexible Image Retrieval Engine) [4] [6]. They are both developed under the GNU license and allow to perform query by example on images, using an image as the starting point for the search process and relying entirely on the image contents.

In the case of GIFT, the complete image database was indexed in a single collection, downscaling each image to 32x32 pixels. For each topic, a visual query is made up of all the images contained in the topic. Then, this visual

query was used in GIFT to obtain the list of the most relevant images (i.e., images which are more similar to those included in the visual query), along with the corresponding relevance values.

On the other hand, we used the results of the FIRE system (kindly provided by the organizers), with no further processing.

2.3. Result Merging

The textual and image result lists are then merged and combined by applying different techniques, which are characterized by an operator and a metric for computing the relevance (score) of the result. **Table 1** shows the defined operators: union (OR), intersection (AND) and external join (LEFT/RIGHT JOIN). Each of these operators selects which images are part of the final result set. Then, results are reranked by computing a new relevance measure value based on their corresponding input results, by using different metrics shown in **Table 2**.

Table 1. Combination operators.

Operators	
OR	$A \cup B$
AND	$A \cap B$
LEFT	$(A \cup B) \cup (A - B)$
RIGHT	$(A \cup B) \cup (B - A)$

Table 2. Score computing metrics.

Metrics	
max	$score = \max(a, b)$
min	$score = \min(a, b)$
avg	$score = \mathrm{avg}(a, b)$
mm	$score = \max(a, b) + \min(a, b) * \dfrac{\min(a, b)}{\max(a, b) + \min(a, b)}$

3. Experiments and Results

Experiments are defined by the choice of different combinations of the previously described modules, operators and score computation metrics. All experiments are fully automatic, avoiding any manual intervention. None of the experiments incorporates relevance feedback, due to time constraints.

We finally submitted a wide set of experiments: 110 multilingual textual (text-based) runs, 22 visual (content-based) runs and 21 mixed runs (using a combination of both). The different name schemas for the run identifiers are described in the following tables: textual retrieval experiments (**Table 3**), visual retrieval experiments (**Table 4**) and mixed experiments (**Table 5**).

Table 3. Textual retrieval experiments.

Run Identifier	Language	Engine	Method	Combination
Txt[Engine][Method]	EN>EN	[1]	[2]	NO
Txt[Combination][Method]	EN>EN	Both	[2]	[3]
Mul[Language][Engine][Method]	X>EN [4]	[1]	[2]	NO
Mul[Language][Combination]	X>EN [4]	Both	NPS	[3]

[1] **Engine:** Xa (Xapian), Lu (Lucene)

[2] **Method:** TI (use image title index), DE (use image description index), NO (use image notes index), LO (use image location index), NP (perform proper noun detection), S (apply stemming)

(3) **Combination:** OR (images in any result set), AND (images in both result sets), max (use maximum value), min (use minimum value), avg (use average value)

(4) **Topic language:** DE (German), ES (Spanish), FR (French), JA (Japanase), PT (Portuguese), RU (Russian), SV (Swedish)

Table 4. Visual retrieval experiments.

Run Identifier	Engine[2]+Combination [2]
Vis[Engine]	[1]
VisG0F0[Combination]	G0+F0 [2]
VisG123[Combination]	G1+G2+G3 [2]
VisG123ANDminF0[Combination]	G123ANDmin+F0 [2]
VisG123ORmaxF0[Combination]	G123ORmax+F0 [2]
VisG123ANDminF0ANDminG0ANDmin	G123ANDmin ANDmin F0 ANDmin G0
VisG123ORmaxF0ORmaxG0ORmax	G123ORmax ORmax F0 ORmax G0

(1) **Engine:** G[1,2,3] (Gift, query with n topic image), G0 (Gift, query with all topic images), F0 (FIRE)

(2) **Combination:** OR (images in any result set), AND (images in both result sets), max (use maximum value), min (use minimum value), avg (use average value)

Table 5. Mixed retrieval experiments.

Run Identifier	Combination
MixLuSG0[Combination]	TxtLuS [1] VisG0
miracleMix1[Combination]	TxtANDminNPS [1] VisG123ANDminF0ANDminG0ANDmin
miracleMix2[Combination]	TxtORmaxNPS [1] VisG123ORmaxF0ORmaxG0ORmax

(1) **Combination:** OR (images in any result set), AND (images in both result sets), LEFT (first OR then AND), max (use maximum value), min (use minimum value), avg (use average value), mm (max-min value)

After the evaluation by the task organizers, results for each kind of experiment are presented in the following tables. Each table shows the run identifier, the mean average precision (MAP), the precision at 10, 20 and 30 first results and the number of relevant images which have been retrieved (out of 3,416 relevant images).

Regarding text-based experiments (**Table 6**), the best results are achieved with the experiment where the title (TI), description (DE) and location (LO) fields of the image caption have been indexed using Xapian (Xa) and applying the proper noun detection module (NP) along with the Porter stemming algorithm (S). For this experiment a MAP of 0.1995 is obtained. Last year, our top ranked text-based experiment had a MAP of 0.2005 with a similar combination of techniques. The only difference is that, while in 2006 topic narratives where also used to build the query, this year there was no narrative provided for the topics. Implications of this are that no improvement has been reached.

Table 6. Results for text-based monolingual experiments.

	MAP	P10	P20	P30	RelRet[1]
TxtXaTIDELONPS	0.1995	0.3033	0.2783	0.2528	1892
TxtORmaxTIDELONPS	0.1972	0.2983	0.2708	0.2461	1892
TxtANDmaxTIDELONPS	0.1960	0.2950	0.2742	0.2456	1832
TxtXaTIDENOLONPS	0.1960	0.2883	0.2683	0.2489	1889
TxtXaNPS	0.1951	0.2933	0.2625	0.2467	1888
TxtORavgTIDELONPS	0.1947	0.2883	0.2592	0.2433	1892
TxtANDavgTIDELONPS	0.1943	0.2850	0.2625	0.2433	1832
TxtORmaxTIDENOLONPS	0.1943	0.2817	0.2617	0.2456	1889
TxtORmaxNPS	0.1936	0.2817	0.2658	0.2506	1888
TxtANDmaxTIDENOLONPS	0.1931	0.2783	0.2658	0.2444	1833
…					
TxtLuTIDELONPS	0.1871	0.2550	0.2400	0.2311	1840

(1) out of 3,416 relevant images

In general, our experiments show Lucene turns out to give worse results than Xapian (the first run with Lucene search engine has a MAP of 0.1871), although this conclusion has to be further studied. Moreover, no strategy for merging results from both search engines have led to any improvement in MAP with respect to the baseline experiments, neither with the OR operator (which was supposed to increase the number of results at the expense of a loss of precision) nor the AND operator (which was supposed to increase precision).

Table 7 shows the precision of the multilingual experiments.

Table 7. Results for text-based multilingual experiments[1].

	Language	MAP	P10	P20	P30	RelRet
MulESXaNPS	Spanish	0.1917	0.2983	0.2633	0.2506	1896
MulESXaTIDELONPS	Spanish	0.1911	0.3033	0.2717	0.2489	1902
MulSVANDmin	Swedish	0.1750	0.2433	0.2000	0.1950	1616
MulSVANDavg	Swedish	0.1726	0.2417	0.2067	0.2000	1616
MulPTORavg	Portuguese	0.1704	0.2433	0.2308	0.2128	1942
MulPTXaTIDELONPS	Portuguese	0.1698	0.2600	0.2425	0.2194	1946
MulRUXaTIDELONPS	Russian	0.1632	0.2600	0.2242	0.2033	1672
MulRUANDavg	Russian	0.1591	0.2550	0.2117	0.1978	1567
MulJALuTIDELONPS	Japanese	0.1545	0.2167	0.1958	0.1878	1518
MulJAORavg	Japanese	0.1532	0.2417	0.2075	0.1794	1533
MulFRXaTIDELONPS	French	0.1419	0.2350	0.2125	0.1950	1771
MulFRXaNPS	French	0.1410	0.2033	0.1950	0.1878	1791
MulDEORavg	German	0.1300	0.1967	0.1683	0.1467	1461
MulDEXaTIDELONPS	German	0.1293	0.1983	0.1733	0.1494	1446

[1] Only the first two results for each language are shown

The best results correspond to Spanish, our mother tongue in which we have a strong expertise, with a MAP similar to the monolingual experiments. We are negatively surprised by the low precision with French and German languages that may be attributed to a deficient stemming module. However, this issue has to be further researched.

The 2007 collection did not include the description field and our experiments have been produced using 2006 data collection (we did not notice this fact until the submission was over). Therefore our results are not really comparable to the results provided by the rest of participants, shown in **Table 8**. However, our baseline experiment would be among the best results this year and our group would be ranked 4[th] out of 16. Our conclusion is that for the IAPR TC-12, textual descriptions of images are very convenient for the retrieval process and should not be dropped out in for the next years.

Table 8. Best text-based experiments at ImageCLEFphoto 2007.

	MAP	P10	P20	P30	RelRet
CINDI_TXT_QE_PHOTO	0.2637	0.4833	0.4092	0.3439	2231
CINDI_TXT_QE_IMG_RF_RERANK	0.2336	0.5150	0.4000	0.3378	2144
EN-EN-AUTO-FB-TXT_FLR	0.2075	0.2650	0.2533	0.2289	2272
budapestacadtextidf	0.2020	0.3033	0.2492	0.2200	2071
INAOE-EN-EN-AN_EXP_5	0.1925	0.3583	0.2942	0.2489	1668

Table 9 in next page shows the results of the visual-based experiments. In this case, MAP is very low due to the complexity of the visual retrieval task in this domain. The best experiment is the combination of results with GIFT and FIRE, and it is better than using GIFT alone. However, we regretfully did not execute the experiment with FIRE alone, which would have been shown whether the combination of both systems improves the final result. Also note that, although the number of relevant images retrieved using AND is lower than when OR is used, the precision is higher.

Table 9. Results for visual-based experiments.

	MAP	P10	P20	P30	RelRet
VisG0F0ANDmin	0.1079	0.3383	0.2400	0.1956	801
VisG123ORmaxF0ANDmin	0.1071	0.3417	0.2450	0.2006	761
VisG0F0ANDmax	0.0933	0.2817	0.1992	0.1667	801
VisG0F0ORmax	0.0916	0.2783	0.1958	0.1606	1053
VisG123ORmaxF0ORmaxG0ORmax	0.0870	0.2700	0.1950	0.1572	1079
VisG0	0.0853	0.2483	0.1867	0.1556	1050
VisG123ORmaxF0ANDmax	0.0853	0.2817	0.1958	0.1628	761
VisG123ORmaxF0ORmax	0.0801	0.2733	0.1808	0.1561	1012
VisG123ANDminF0ORmax	0.0797	0.2067	0.1600	0.1439	1263
VisG123ORmax	0.0785	0.2683	0.1800	0.1544	1009
VisG123ANDminF0ORmin	0.0778	0.2117	0.1783	0.1450	1262

Results from other participants, shown in **Table 10**, clearly outperform our own results. Our MAP is 56% lower than the MAP of the best experiment and our group is ranked 7[th] out of 13 participants.

Table 10. Best visual-based experiments at ImageCLEFphoto 2007.

	MAP	P10	P20	P30	RelRet
INAOE-VISUAL-EN-AN_EXP_3	0.1925	0.3583	0.2942	0.2489	1668
AUTO-NOFB-IMG_COMBFK	0.1890	0.4700	0.3517	0.2922	1708
IMG-dcu_uta-dyn-equal-fire	0.1340	0.3600	0.2658	0.2217	1339
IPAL04V_12RUNS_WEIGHT	0.1204	0.3367	0.2525	0.2083	1330
INAOE-LF_CBR	0.1200	0.2267	0.1908	0.1689	2000

Table 11 shows the results for mixed (textual and visual) experiments.

Table 11. Results for mixed textual and visual retrieval experiments.

	MAP	P10	P20	P30	RelRet
Mix2LEFTmm	0.2244	0.4450	0.3617	0.3067	1888
MixLuSG0LEFTmm	0.2027	0.4200	0.3517	0.2872	1812
Mix2LEFTmax	0.1985	0.3650	0.3142	0.2739	1888
Mix1LEFTmm	0.1975	0.3067	0.2742	0.2500	1763
Mix1LEFTmax	0.1915	0.2900	0.2642	0.2433	1763
Mix2LEFTavg	0.1901	0.2717	0.2592	0.2489	1888
Mix2LEFTmin	0.1901	0.2717	0.2592	0.2489	1888
Mix1LEFTavg	0.1888	0.2733	0.2575	0.2411	1763

The best experiment uses a combination of all textual and visual results obtained in the previous experiments with the LEFT operator along with the Max-Min (mm) scoring metric. Moreover, the mm metric clearly outperforms the others (max, min, avg) even with the same operator over the same result sets, as seen in Mix2LEFT experiments. The most interesting conclusion is that this merging strategy is successful as the MAP is higher than in both the textual and visual partial experiments with which the experiment is built.

Finally, although our mixed experiments can really be compared to the others groups' for the reasons already stated before, our MAP is considerably lower than the MAP of the best mixed experiment (**Table 12**), with a decrease to 70%. However, MIRACLE ranked 5[th] (out of over 10 participants), which is indeed considered to be an acceptable position.

Table 12. Best mixed experiments at ImageCLEFphoto 2007.

	MAP	P10	P20	P30	RelRet
cut-EN2EN-F50	0.3175	0.5900	0.4592	0.3839	2251
EN-EN-AUTO-FB-TXTIMG_MPRF	0.3168	0.4900	0.4167	0.3711	2581
EN-EN-AUTO-FB-XTIMG_QTXT_COMBPREFFKTXT	0.3020	0.4300	0.3733	0.3306	2595
DE-EN-AUTO-FB-TXTIMG_MPRF	0.2899	0.4483	0.3883	0.3528	2572
cut-EN2EN-F20	0.2846	0.5283	0.4025	0.3433	2191

4. Conclusions and Future Work

Results in general show a poor performance for all groups, due to the characteristics of the image collection and the difficulty of the defined topics.

For the textual retrieval, the experiment where the title, description and location fields of the image caption have been indexed using Xapian and applying the proper noun detection module and the stemming component has produced the best results, with a MAP of 0.1995, which shows that no improvement has been reached with regards to last year experiments when our top ranked text-based experiment had a MAP of 0.2005 with a similar combination of techniques. In general, we obtain worse MAP with Lucene rather than Xapian, although this topic has to be further investigated. Regarding the visual retrieval, our results are not good compared to other groups', but it is not strange as our areas of expertise do not specifically include image processing and we resorted to publicly available "black-box" engines. In addition, no specific conclusion can be obtained as we omitted to include the result for FIRE itself, so we could not evaluate if the merging strategy is successful or not.

However, our main interest was not in experiments where only text or image content is used in the retrieval process. Instead, our challenge was to test whether the text-based image retrieval could improve the analysis of the content of the image, or vice versa. The most interesting conclusion is that merging strategies are successful as our best mixed experiment outperforms both the textual and visual experiments in which it is based, using the LEFT operator for the combination along with the Max-Min (mm) metric for computing the relevance. This result shows that our hypothesis was right. Our combination of a "black-box" search using publicly accessible content-based retrieval engines with a text-based search has turned out to provide better results than other presumably "more complex" techniques. This simplicity may be a good starting point for the implementation of a real system.

We are sure that there may still be some place for improvement with a more careful study of the parameters for the merging strategy (both the combination operator and the score computing metric). In addition, none of the experiments incorporates relevance feedback, due to time constraints. However, we are already carrying out some experiments that incorporate this technique which are showing promising results, overcoming our submissions.

Acknowledgements

References

[1] Apache Lucene project. On line http://lucene.apache.org [Visited 10/08/2007].

[2] Charniak, Eugene. A Maximum-Entropy-Inspired Parser. In Proceedings of NAACL-2000, 2000.

[3] CLEF 2005 Multilingual Information Retrieval resources page. On line http://www.computing.dcu.ie/~gjones/CLEF2005/Multi-8/ [Visited 10/08/2007].

[4] Deselaers, T.; Keysers; D.; Ney, H. FIRE - Flexible Image Retrieval Engine: ImageCLEF 2004 Evaluation. In CLEF 2004, LNCS 3491, Bath, UK, pp 688-698, September 2004.

[5] Eurowordnet: Building a Multilingual Database with Wordnets for several European Languages. March (1996). On line http://www.illc.uva.nl/EuroWordNet/ [Visited 10/08/2007].

[6] FIRE: Flexible Image Retrieval System. On line http://www-i6.informatik.rwth-aachen.de/~deselaers/fire.html [Visited 10/08/2007].

[7] GIFT: The GNU Image-Finding Tool. On line http://www.gnu.org/software/gift/ [Visited 10/08/2007].

[8] Grubinger, Michael; Clough, Paul; Hanbury, Allan; Müller, Henning. Overview of the ImageCLEF 2007 Photographic Retrieval Task. Working Notes of the 2007 CLEF Workshop, Budapest, Hungary, September 2007.

[9] Grubinger, Michael; Clough, Paul; Müller, Henning; Deselaers, Thomas. The IAPR-TC12 benchmark: A new evaluation resource for visual information systems. In International Workshop OntoImage'2006 Language Resources for Content-Based Image Retrieval, held in conjunction with LREC'06, pages 13–23, Genoa, Italy, May 2006.

[10] IAPR: On line http://www.iapr.org/ [Visited 10/08/2007]

[11] Martínez-Fernández, J.L.; Villena-Román, Julio; García-Serrano, Ana M.; Martínez-Fernández, Paloma. MIRACLE team report for ImageCLEF IR in CLEF 2006. Proceedings of the Cross Language Evaluation Forum 2006, Alicante, Spain. 20-22 September 2006.

[12] Martínez-Fernández, J.L.; Villena-Román, Julio; García-Serrano, Ana M.; González-Cristóbal, José Carlos. Combining Textual and Visual Features for Image Retrieval. Accessing Multilingual Information Repositories: 6th Workshop of the Cross-Language Evaluation Forum, CLEF 2005, Vienna, Austria, Revised Selected Papers. Carol Peters et al (Eds.). Lecture Notes in Computer Science, Vol. 4022, 2006. ISSN: 0302-9743.

[13] Montoyo, A. Método basado en marcas de especificidad para WSD. In Proceedings of SEPLN (Sociedad Española para el Procesamiento del Lenguaje Natural), nº 24, September 2000.

[14] Porter, Martin. Snowball stemmers and resources page. On line http://www.snowball.tartarus.org [Visited 10/08/2007].

[15] University of Neuchatel. Page of resources for CLEF (Stopwords, transliteration, stemmers …). On line http://www.unine.ch/info/clef [Visited 10/08/2007].

[16] Xapian: an Open Source Probabilistic Information Retrieval library. On line http://www.xapian.org [Visited 10/08/2007].